

Oden Institute REPORT 20-14

July 2020

Matrix-free isogeometric Galerkin method for Karhunen-Loève approximation of random fields using tensor product splines, tensor contraction and interpolation based quadrature

by

Michal L. Mika, Thomas J.R. Hughes, Dominik Schillinger, Peter Wriggers, and René R. Hiemstra



Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
Austin, Texas 78712

Reference: Michal L. Mika, Thomas J.R. Hughes, Dominik Schillinger, Peter Wriggers, and René R. Hiemstra, "Matrix-free isogeometric Galerkin method for Karhunen-Loève approximation of random fields using tensor product splines, tensor contraction and interpolation based quadrature," Oden Institute REPORT 20-14, Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, July 2020.

Oden Institute REPORT 20-14

July 2020

Matrix-free isogeometric Galerkin method for Karhunen-Loève approximation of random fields using tensor product splines, tensor contraction and interpolation based quadrature

by

Michai L. Mika, Thomas J.R. Hughes, Dominik Schillinger, Peter Wriggers, and René R. Hiemstra



Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
Austin, Texas 78712

Reference: Michai L. Mika, Thomas J.R. Hughes, Dominik Schillinger, Peter Wriggers, and René R. Hiemstra, "Matrix-free isogeometric Galerkin method for Karhunen-Loève approximation of random fields using tensor product splines, tensor contraction and interpolation based quadrature," Oden Institute REPORT 20-14, Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, July 2020.

Matrix-free isogeometric Galerkin method for Karhunen-Loève approximation of random fields using tensor product splines, tensor contraction and interpolation based quadrature

Michał L. Mika^{a,*}, Thomas J.R. Hughes^b, Dominik Schillinger^a, Peter Wriggers^c, René R. Hiemstra^a

^a*Institut für Baumechanik und Numerische Mechanik, Leibniz Universität Hannover*

^b*Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin*

^c*Institut für Kontinuumsmechanik, Leibniz Universität Hannover*

Abstract

The Karhunen-Loève series expansion decomposes a stochastic process into an infinite series of pairwise uncorrelated random variables and pairwise L^2 -orthogonal functions. For any given truncation order of the infinite series the basis is optimal in the sense that the total mean squared error is minimized. The orthogonal basis functions are determined as the solution of an eigenvalue problem corresponding to the homogeneous Fredholm integral equation of the second kind, which is computationally challenging for several reasons. Firstly, a Galerkin discretization requires numerical integration over a $2d$ dimensional domain, where d , in this work, denotes the spatial dimension. Secondly, the main system matrix of the discretized weak-form is dense. Consequently, the computational complexity of classical finite element formation and assembly procedures as well as memory requirements of direct solution techniques become quickly computationally intractable with increasing polynomial degree, number of elements and degrees of freedom. In this work we present a matrix-free solution strategy that scales favorably with problem size and polynomial degree and is embarrassingly parallel. Our approach is based on (1) an interpolation based quadrature that minimizes the required number of quadrature points; (2) inexpensive reformulation of the generalized eigenvalue problem into a standard eigenvalue problem; and (3) a matrix-free and multithreaded matrix-vector product for iterative eigenvalue solvers. Two higher-order three-dimensional benchmarks illustrate exceptional computational performance combined with high accuracy and robustness.

Keywords:

Matrix-free solver, Kronecker products, random fields, Fredholm integral eigenvalue problem, isogeometric analysis

Contents

| | | |
|----------|---|----------|
| 1 | Introduction | 2 |
| 1.1 | Discrete representation of random fields by the truncated Karhunen-Loève series expansion | 3 |
| 1.2 | Challenges in numerical solution of the KLE by means of the Galerkin method | 3 |
| 1.3 | Splines as a basis for random fields | 4 |
| 1.4 | Contributions | 5 |
| 1.5 | Outline | 6 |

*Corresponding author

Email addresses: mika@ibnm.uni-hannover.de (Michał L. Mika), hughes@ices.utexas.edu (Thomas J.R. Hughes), schillinger@ibnm.uni-hannover.de (Dominik Schillinger), wriggers@ikm.uni-hannover.de (Peter Wriggers), rene.hiemstra@ibnm.uni-hannover.de (René R. Hiemstra)

| | | |
|----------|--|-----------|
| 2 | Background and notation | 6 |
| 2.1 | Evaluation of the computational cost of an algorithm | 6 |
| 2.2 | Kronecker products and tensor contraction | 6 |
| 2.3 | B-splines | 7 |
| 2.4 | Geometrical mapping | 7 |
| 3 | Isogeometric Galerkin discretization of the Karhunen-Loève series expansion | 8 |
| 3.1 | KL expansion of random fields | 8 |
| 3.2 | Truncated KL expansion | 10 |
| 3.3 | Variational formulation | 10 |
| 3.4 | Choice of the trial space | 10 |
| 3.5 | Matrix formulation | 11 |
| 4 | Efficient matrix-free solution strategy | 12 |
| 4.1 | Reformulation into a standard algebraic eigenvalue problem | 12 |
| 4.2 | Interpolation based quadrature for integral equations | 13 |
| 4.3 | Matrix formation | 14 |
| 4.4 | Matrix-free solution strategy | 15 |
| 4.4.1 | Basic setup | 15 |
| 4.4.2 | Matrix-free algorithm | 15 |
| 5 | Computational complexity analysis | 16 |
| 5.1 | Standard finite element procedures | 16 |
| 5.2 | Finite element procedures employing sum factorization | 17 |
| 5.3 | Proposed strategy using interpolation based quadrature | 17 |
| 5.4 | Storage comparison | 18 |
| 6 | Numerical results | 18 |
| 6.1 | One-dimensional case study | 18 |
| 6.2 | Random field with exponential kernel in a three-dimensional half-open cylindrical domain . . | 19 |
| 6.3 | Random field with Gaussian kernel in a three-dimensional hemispherical shell | 24 |
| 7 | Conclusion | 28 |

1. Introduction

Most physical systems exhibit randomness, which, because of its lack of pattern or regularity, can not be explicitly captured by deterministic mathematical models. The randomness may be due to the nature of the phenomenon itself, called *aleatoric* uncertainty, or due to a lack of knowledge about the system, referred to as *epistemic* uncertainty. In the latter the uncertainty may be reduced by obtaining additional data about the system at hand. An example of an epistemic uncertainty encountered in engineering are the fluctuations of material properties throughout a body, which occur due to the inhomogeneity of the medium. Deterministic mechanical models typically feature empirically derived material parameters, such as material stiffness and yield stress, that are assumed constant throughout the body. Their value is typically determined as a statistical volumetric average over a large set of laboratory specimens. This idealized model of reality may be insufficient in e.g. structural risk or reliability analysis and prediction, which is concerned with probabilities of violation of safety limits or performance measures, respectively [40]. In this case the effects of uncertainty on the result of a computation need to be quantified.

Uncertainty in physical quantities that vary in space and or time may be adequately modeled by *stochastic processes* or *random fields* [54]. This approach generalizes a deterministic system modeled by a partial differential equation to a stochastic system modeled by a stochastic partial differential equation or SPDE. Reliable predictions may be obtained by propagating uncertainties in input variables to those in the response.

The main objective is to compute the response statistics, such as the mean and variance in the random solution field, or the probability that a set tolerance is exceeded. To compute these statistics it is necessary to discretize the SPDE, not only in space and time, but also in the stochastic dimensions. This can be a complicated task, not because of modeling randomness, but due to a *curse of dimensionality*. Every random variable contributes one dimension to the problem. Hence, it is important to keep their total to a minimum.

1.1. Discrete representation of random fields by the truncated Karhunen-Loève series expansion

One of the relevant questions in stochastic analysis is how to represent random fields discretely, suitable for use in numeric computation. The essential step is to break down the representation into a tractable number of mutually independent random variables, whose combination preserves the stochastic variability of the process [13, 27]. One representation that is of particular interest is the truncated *Karhunen-Loève series expansion* or *KLE* [26, 38]. The KLE decomposes a stochastic process into an infinite series of pairwise uncorrelated random variables and pairwise L^2 -orthogonal basis functions¹. Truncating the series expansion after M terms yields the best M -term linear approximation of the random field, in the sense that the total mean squared error is minimized.

Computation of the truncated KL expansion requires the solution of a homogeneous Fredholm integral eigenvalue problem (IEVP) of the second kind. In general this is only possible numerically. The most popular numerical methods to solve IEVP's are the Nyström method, degenerate kernel methods and the collocation and Galerkin method [4, 31]. The Galerkin method is widely regarded as superior due to its approximation properties and solid theoretical foundation [13]. Within the trial space under consideration, the Galerkin method produces the best L^2 fit of the analytical modes. Furthermore, the discrete modes preserve exactly the L^2 orthogonality property of the analytical mode-shapes.

1.2. Challenges in numerical solution of the KLE by means of the Galerkin method

Efficient solution of the KLE using the Galerkin method is a computationally challenging task [13]. The main challenges are the following:

- (i) A Galerkin discretization requires numerical integration over a $2d$ dimensional domain, where d , in this work, denotes the spatial dimension. The computational complexity of classical finite element formation and assembly procedures scales as $\mathcal{O}(N_e^2 \cdot p^{3d})$, where N_e is the global number of elements, p the polynomial degree and d the spatial dimension.
- (ii) The main system matrix of the discretized weak-form is dense and requires $8N^2$ bytes of memory in double precision arithmetic, where N is the dimension of the trial space.
- (iii) Numerical solution requires one sparse backsolve $\mathcal{O}(N^2)$ and one dense matrix-vector product $\mathcal{O}(N^2)$ in each iteration of the eigenvalue solver, thus the solution time of the numerical eigenvalue solver scales $\mathcal{O}(N^2 \cdot N_{\text{iter}})$, where N_{iter} is the number of iterations required by the Lanczos solver.

Table 1 illustrates that explicit storage of the dense system matrix requires impracticable amounts of memory for problems involving more than $100K$ degrees of freedom. Hence, the computational complexity of classical finite element formation and assembly procedures as well as memory requirements of direct solution techniques become quickly computationally intractable with increasing polynomial degree, number of elements and degrees of freedom.

There has been a particular research effort devoted to alleviating the disadvantages of the Galerkin method. In [1, 16, 29] an approximation by (Kronecker product) hierarchical matrices is used to efficiently compute the dense matrices, as well as to reduce the memory requirements. These matrices are sparse and allow for matrix multiplication, addition and inversion in $\mathcal{O}(N \log N)$ time (or for Kronecker product hierarchical matrices in $\mathcal{O}(N)$ time) where N is the number of degrees of freedom. The generalized Fast Multipole Method, which also scales with $\mathcal{O}(N \log N)$, has been proposed in [52]. This method was shown

¹In analogy with a Fourier series expansion we also refer to the basis functions as KL modes or mode-shapes.

to *not* yield significant speed-ups for p finite element methods and thus it is recommended for kernels of low regularity. Wavelet Galerkin-schemes [44] are also being used and can be coupled with compression techniques for boundary value problems [9], but have the disadvantage, that the number of eigenmodes to be computed must be known in advance. The pivoted Cholesky decomposition [17] focuses on approximating the discretized random fields with sufficiently fast decaying eigenvalues. In this case a truncation of the pivoted Cholesky decomposition of the covariance operator allows for an estimation of the eigenvalues in the post-processing step in $\mathcal{O}(M^2N)$ time, where M is the truncation order of the Cholesky decomposition. One of the advantages of this method is the fact, that the number of eigenmodes required for a certain accuracy of the random field discretization can be estimated in advance.

Table 1: Minimum memory required for storage of the main system matrix in the solution of the homogeneous Fredholm integral problem of the second kind assuming double-precision floating point arithmetic.

| Number of degrees of freedom | 10^3 | 10^4 | 10^5 | 10^6 |
|------------------------------|--------|--------|--------|--------|
| Matrix storage | 8 MB | 800 MB | 80 GB | 8 TB |

1.3. Splines as a basis for random fields

Splines are piecewise polynomials with increased smoothness across element boundaries compared to classical finite elements. Traditionally, splines have been primarily used as shape functions in computer aided design. More recently, with the introduction of isogeometric analysis [21], splines have become more established as trial functions in finite element analysis. Although isogeometric analysis was originally introduced to improve the interoperability across several stages of the design to analysis process, it has proven its fidelity as an analysis technology. We refer to the monograph [7] and references contained therein for an exposition of isogeometric analysis applied to deterministic problems in structural and fluid mechanics.

More recently, spline based isogeometric analysis has found its way into the stochastic community. Stochastic methods have been proposed to quantify uncertainty due to material randomness in linear elasticity [34, 25], static analysis of plates [58], vibrational analysis of shells [37], static and dynamic structural analysis of random composite structures [11] and functionally graded plates [20, 35, 36]. In [60] a method is proposed to quantify the effect due to uncertainty in shape. Of these, the methods proposed in [34, 35, 36, 37] use isogeometric analysis within a spectral stochastic finite element framework [14], which is based on a KL expansion of random fields. The methods in [11, 20, 58] use perturbation series of which [58] expands random fields in terms of the KLE. Standard polynomial chaos is used in [60], while the methods in [12] and [25] propose new chaos expansions in terms of tensor product B-spline basis functions; thus extending generalized polynomial chaos [59] to spline chaos.

To the best of our knowledge, [2] is the first work in which splines have been used to approximate the truncated KLE. In his work the author applies a degenerate kernel approximation based on tensor product spline interpolation at the Greville abscissa. More recently, in the spirit of isogeometric analysis, non-uniform rational B-splines (NURBS) have been used to approximate the KLE using the Galerkin method [47] and the collocation method [24]. These methods avoid the geometrical errors in the representation of CAD geometry typically made within the classical finite element method. The authors note that the use of splines in the geometry description as well as in discretization of the spatial and stochastic dimensions could enable a “seamless uncertainty quantification pipeline”.

In the context of the present work we would like to highlight the superior spectral approximation properties of smooth splines as compared to classical C^0 finite element shape functions. Several studies [6, 22, 23, 46] have investigated the spectral approximation properties of splines in eigenvalue problems corresponding to second and fourth order differential operators and have demonstrated that splines have improved robustness and accuracy per degree of freedom across virtually the entire range of modes. The numerical results for the Fredholm integral eigenvalue problem are no different, as corroborated by the results shown in Figure 1. It’s precisely these properties that make splines appealing in the representation of random fields by means of the Karhunen-Loève expansion.




Figure 1: Normalized discrete eigenvalues corresponding to a univariate Fredholm integral eigenvalue problem with an exponential kernel (correlation length is one). Comparison of eigenvalues obtained by C^1 quadratic splines to C^0 quadratic piecewise polynomials. Both methods employ a standard Galerkin projection based on full Gauss quadrature. The reference solution used to normalize the results is computed by the approach described in Remark 6.2

1.4. Contributions

We present a matrix-free isogeometric Galerkin method for Karhunen-Loève approximation of random fields by splines. Our solution methodology resolves several of the aforementioned computational bottlenecks associated with numerical solution of integral eigenvalue problems and enables solution of large-scale three-dimensional IEVPs on complex domains. Below we summarize our main contributions.

Conversion to a standard eigenvalue problem

We have chosen a specific trial space of rational spline functions whose Gramian matrix has a Kronecker product structure independent of the geometric mapping. This enables us to perform the backsolve, used to convert the IEVP to standard form, in $\mathcal{O}(N \cdot N^{1/d})$ time by utilizing standard linear algebra techniques from [15, 49].

Interpolation based quadrature

We present an interpolation based quadrature technique designed and optimized specifically for the variational formulation of the Fredholm integral equation. The approach integrates a rich target space of functions with minimal number of quadrature points and outperforms existing competitive techniques in isogeometric analysis, such as quadrature by interpolation and table look-up [39, 43] and weighted quadrature [5, 19, 51]. The proposed interpolation based quadrature technique is inspired by a similar technique used within linear finite elements in [28, Chapter 3.1.3] and [29]. Instead, our approximation of the covariance function is based on higher order tensor product spline interpolation and resembles the kernel approximation made in [2]. Besides requiring as few quadrature points as possible, the interpolation based quadrature technique exposes Kronecker structure in the integral equations, reducing computational complexity significantly.

Matrix-free solution methodology

We present a matrix-free solution methodology to avoid explicit storage of the dense system matrix associated with numerical computation of the KLE. The matrix-free solution methodology not only reduces

the memory complexity from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$, but also significantly reduces the solution time. This is achieved by integrating the matrix-free solver with the proposed interpolation based quadrature technique. The latter exposes Kronecker structure in the resulting discrete integral equation, thereby reducing formation costs to $\mathcal{O}(N \cdot N^{1/d})$ per iteration, leaving only the dense matrix-vector product that is associated with the Lanczos algorithm that remains $\mathcal{O}(N^2)$ per iteration. The integrative approach to quadrature and matrix-free solution techniques, exploiting Kronecker structure, is inspired by the matrix-free weighted quadrature method proposed recently in [51].

Open source implementation

We provide an open-source Python implementation of the described techniques that is available for download at <https://github.com/mika05/tensiga>.

1.5. Outline

In Section 2 we briefly review the necessary mathematical and algorithmic background with regard to Kronecker products, B-splines and NURBS. In Section 3 we present the Karhunen-Loève series expansion of random fields and the weak formulation of the corresponding Fredholm integral eigenvalue problem of the second kind. In Section 4 we introduce our methodology for numerical solution of the truncated KLE. This includes reformulation of the eigenvalue problem to standard form, interpolation based quadrature of the weak form of the Fredholm integral problem, and a matrix-free algorithm with low computational complexity and minimal memory requirements that is embarrassingly parallelizable. The computational complexity is described in more detail in Section 5, where we compare our method with usual formation and assembly techniques used for standard Galerkin methods from the literature. Finally, in Section 6, we present a one-dimensional numerical study and several three-dimensional high-order numerical examples. A conclusion and an outlook with recommendations for future work are given in Section 7.

2. Background and notation

This section introduces some of the machinery that is used throughout the paper. The presented solution methodology for the Fredholm integral equation relies heavily on the properties of Kronecker products in combination with multidimensional tensor contraction [15]. We briefly review the main properties used in this work and illustrate their use in the Kronecker matrix-vector product. The Kronecker structure of the involved matrices is a direct consequence of the chosen tensor product spline function spaces. We briefly introduce B-splines as a basis for polynomial splines and Non-Uniform Rational B-splines (NURBS) for smooth geometrical mappings. For additional details we refer the reader to standard reference books [7, 45].

2.1. Evaluation of the computational cost of an algorithm

The computational cost of the algorithms discussed in this work are evaluated in terms of *floating point operations per second (flops)*. A single *flop* represents the amount of work required to perform one floating point addition, subtraction, multiplication or division [15]. Although the number of flops does not provide a complete assessment of the efficiency of an algorithm, it is widely used in the literature. Indeed, many other considerations such as cache-line efficiency and number of memory allocations can have a large impact on the performance of an algorithm. Typically, we are interested in the leading terms that dominate the computational cost of an algorithm and record the performance in terms of an order-of-magnitude estimate of the number of flops, written in Big-Oh notation as $\mathcal{O}(\cdot)$.

2.2. Kronecker products and tensor contraction

Let $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{p \times q}$ and $C \in \mathbb{R}^{s \times t}$ denote real valued matrices. The Kronecker product $A \otimes B \in \mathbb{R}^{m \cdot p \times n \cdot q}$ is a matrix defined as

$$A \otimes B := \begin{bmatrix} A_{11}B & \cdots & A_{1n}B \\ \vdots & & \vdots \\ A_{m1}B & \cdots & A_{mn}B \end{bmatrix} \quad (1)$$

Kronecker products satisfy the following properties

$$(A \otimes B) \otimes C = A \otimes (B \otimes C) \quad (\text{associativity}) \quad (2a)$$

$$(A \otimes B)(C \otimes D) = (AC) \otimes (BD) \quad (\text{mixed product property}) \quad (2b)$$

$$(A \otimes B)^{-1} = A^{-1} \otimes B^{-1} \quad (\text{inverse of a Kronecker product}) \quad (2c)$$

$$(A \otimes B)^\top = A^\top \otimes B^\top \quad (\text{transpose of a Kronecker product}) \quad (2d)$$

Let $X \in \mathbb{R}^{n_1 \times \dots \times n_d}$, $Y \in \mathbb{R}^{m_1 \times \dots \times m_d}$ denote two d -dimensional arrays. Vectorization of X is a linear operation that maps X to a vector $\text{vec}(X) \in \mathbb{R}^{n_1 \dots n_d}$ with entries

$$\text{vec}(X)_i := X_{i_1 \dots i_d}, \quad \text{where } i = i_1 + (i_2 - 1)n_1 + (i_3 - 1)n_1 \cdot n_2 + \dots + (i_d - 1)n_1 \cdot \dots \cdot n_{d-1}. \quad (3)$$

One recurring theme in this paper involving Kronecker matrices is efficient matrix vector multiplication. Let $D_k \in \mathbb{R}^{m_k \times n_k}$ denote a set of d matrices $\{D_{i_k j_k}, k = 1, \dots, d\}$, $i_k = 1, \dots, m_k$ and $j_k = 1, \dots, n_k$. The matrix-vector product

$$\text{vec}(Y) = (D_d \otimes \dots \otimes D_1) \text{vec}(X) \quad \mathcal{O}(M \cdot N) \text{ flops} \quad (4a)$$

can be written as a tensor contraction instead

$$Y_{i_1 \dots i_d} = \sum_{j_1 \dots j_d} D_{i_1 j_1} \dots D_{i_d j_d} X_{j_1 \dots j_d} \quad \mathcal{O}(\max(N \cdot m_1, n_d \cdot M)) \text{ flops} \quad (4b)$$

Here $N = n_1 \dots n_d$ and $M = m_1 \dots m_d$. The second approach scales nearly linearly with matrix size and significantly outperforms standard matrix vector multiplication which scales quadratically with the matrix size. In practice, highly optimized linear tensor algebra libraries can be used to perform the tensor contraction such as the tensor algebra compiler (TACO) [30]. Our Python implementation uses Numpy's reshaping and matrix-matrix product routines, which call low-level BLAS routines. The implemented reshapes do not require any expensive and unnecessary data copies.

2.3. B-splines

Consider a d -dimensional parametric domain $\hat{\mathcal{D}} = [0, 1]^d \subset \mathbb{R}^d$ with local coordinates $\hat{x} = (\hat{x}_1, \dots, \hat{x}_d)$. Let $(B_{i_k, p_k}(\hat{x}_k), i_k = 1, \dots, n_k)$ denote the univariate B-spline basis of polynomial degree p_k and dimension n_k , corresponding to the k th parametric coordinate \hat{x}_k . We consider multivariate B-splines as tensor products of univariate B-splines

$$B_i(\hat{x}) = \prod_{k=1}^d B_{i_k, p_k}(\hat{x}_k), \quad i := (i_1, \dots, i_d). \quad (5)$$

Here $i \in \mathcal{I}$ is a multi-index in the set $\mathcal{I} := \{(i_1, \dots, i_d) : 1 \leq i_k \leq n_k\}$. The collection of all multivariate B-spline basis functions spans the space

$$\mathcal{B}_h := \text{span}\{B_i(\hat{x})\}_{i \in \mathcal{I}}. \quad (6)$$

It is important to note that splines allow for increased continuity between polynomial elements as compared to classical C^0 -continuous finite elements basis functions. This turns out to have significant impact on the spectral accuracy of the Galerkin method. This is evidenced by several studies [6, 22, 23, 46] and will be discussed in some detail in this work.

2.4. Geometrical mapping

Let $F : \hat{\mathcal{D}} \rightarrow \mathcal{D}$ map a point \hat{x} from the parametric domain $\hat{\mathcal{D}}$ to a point x in the physical domain \mathcal{D} . We assume that the map F and its inverse are smooth such that the Jacobian matrix $[DF(\hat{x})]_{ij} := \frac{\partial F_i}{\partial \hat{x}_j}$

and its inverse are well-defined. In this work F is represented as a linear combination of Non-Uniform Rational B-splines (NURBS). NURBS are rational functions of B-splines that enable representation of common engineering shapes with conic sections, which cannot be represented by polynomial B-splines [45]. The discretization method presented in this work makes heavy use of tensor product properties of the involved function spaces. Since NURBS do not have a tensor product structure, we use them only to represent the geometry and do not consider them as a basis for the function spaces.

3. Isogeometric Galerkin discretization of the Karhunen-Loève series expansion

The Karhunen-Loève (KL) series expansion decomposes a stochastic process or field into an infinite linear combination of L^2 -orthogonal functions and uncorrelated stochastic random variables. In this section we present the probability theory underlying the KL expansion and discuss its discretization by means of the Galerkin method.

3.1. KL expansion of random fields

Consider a complete probability space $(\Theta, \Sigma, \mathbb{P})$. Here Θ denotes a sample set of random events, Σ is the σ -algebra of Borel subsets of Θ and \mathbb{P} is a probability measure $\mathbb{P} : \Sigma \rightarrow [0, 1]$. A random field $\alpha(\cdot, \theta) : \Theta \mapsto L^2(\mathcal{D})$ on a bounded domain $\mathcal{D} \in \mathbb{R}^d$ is a collection of deterministic functions of $x \in \mathcal{D}$, called realizations, that are indexed by events $\theta \in \Theta$. A subset of realizations $\alpha(\cdot, \Theta_s)$, $\Theta_s \in \Sigma$, has a probability of occurrence of $\mathbb{P}(\Theta_s)$.

Let $\mathbb{E}[\cdot]$ denote the expectation operator corresponding to the probability measure \mathbb{P} . Assuming $\alpha \in L^2(\mathcal{D} \times \Theta)$ its first and second order moments exist and are given by

$$\mu(x) := \mathbb{E}[\alpha(x, \theta)] \quad \text{and} \quad (7a)$$

$$\Gamma(x, x') := \mathbb{E}[(\alpha(x, \theta) - \mu(x))(\alpha(x', \theta) - \mu(x'))] \quad (7b)$$

Here μ is called the mean or expected value of α over all possible realizations, and $\Gamma : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ is called its *covariance function* or *kernel*. By definition, the kernel is bounded, symmetric and positive semi-definite [14]. Because the kernel is square integrable, that is, $\Gamma \in L^2(\mathcal{D} \times \mathcal{D})$, it is in fact a Hilbert-Schmidt kernel.

A random field is *stationary* or *homogeneous* if its statistical properties do not vary as a function of position $x \in \mathcal{D}$. This implies that the covariance function can be written as a function of the difference $x - x'$. Furthermore, for *isotropic* random fields the statistical properties are invariant under rotations, which means the covariance is a function of Euclidean distance $\|x - x'\|_2$.

Remark 3.1. *Although, the Euclidean distance is widely used in the literature its use is not always justified. In general, the geodesic distance, i.e. the shortest distance between points x and x' along all paths contained in \mathcal{D} , is the true measure of distance. The Euclidean distance can vary significantly from the geodesic distance especially if the correlation length is relatively large and the domain is non-convex. The geodesic distance is, however, difficult and expensive to compute, which explains its non-use. In this work we also use the Euclidean distance measure and assume its choice is a reasonable one in the context of the applied numerical benchmark problems.*

Figure 2 shows two common examples of covariance functions that correspond to stationary isotropic random fields: the *exponential* and the *Gaussian* or *squared exponential kernel*. Important parameters that influence the locality of these correlation functions are the variance σ^2 and correlation length bL . Here L denotes a characteristic length and b is a dimensionless factor.

The KL expansion of a random field $\alpha(\cdot, \theta)$ requires the solution of an integral eigenvalue problem. Consider the linear operator

$$T : L^2(\mathcal{D}) \mapsto L^2(\mathcal{D}), \quad (T\phi)(x) := \int_{\mathcal{D}} \Gamma(x, x')\phi(x') \, dx'. \quad (8)$$

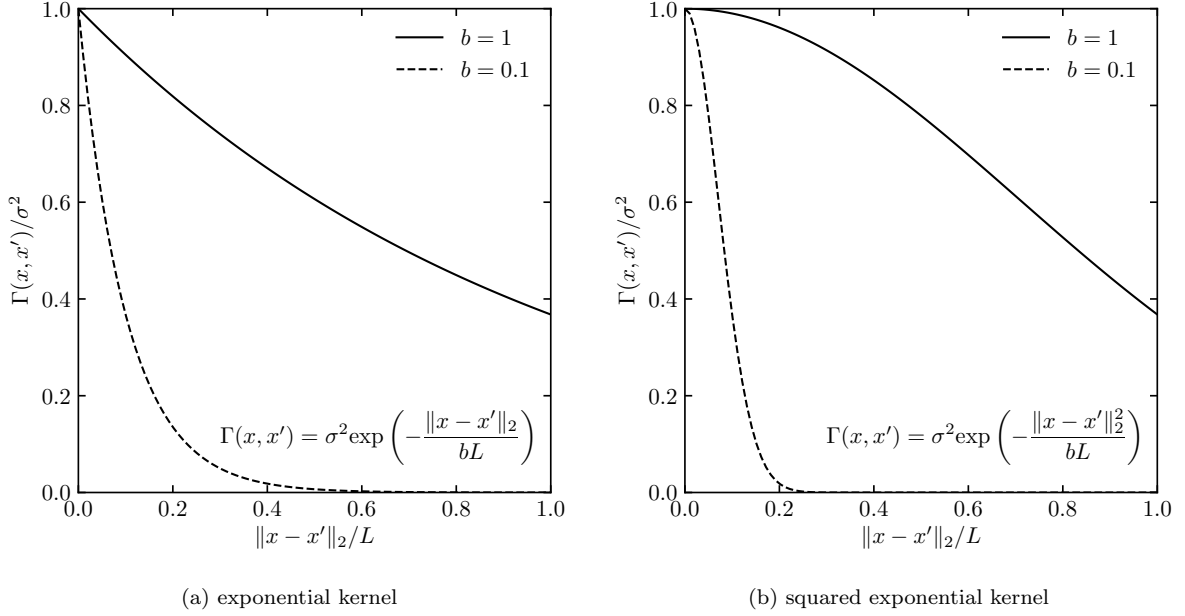


Figure 2: The exponential and squared exponential (Gaussian) covariance functions for different correlation lengths with $b = \{0.1, 1.0\}$. Note the difference in the continuity of both kernels at $x = x'$. The exponential kernel is C^0 , while the square exponential kernel is C^∞ at $x = x'$.

The operator T is compact. In fact, T is a Hilbert-Schmidt operator, since the covariance function is a Hilbert-Schmidt kernel. Furthermore, since the covariance function is positive semi-definite and symmetric [14], T is a self-adjoint positive semi-definite linear operator. The eigenfunctions $\{\phi_i\}_{i \in \mathbb{N}}$ of T are defined by the homogeneous Fredholm integral eigenvalue problem of the second kind,

$$T\phi_i = \lambda_i \phi_i, \quad \phi_i \in L^2(\mathcal{D}) \text{ for } i \in \mathbb{N}. \quad (9)$$

Important properties of the eigenpairs are (1) the normalized eigenfunctions $\{\phi_i\}_{i \in \mathbb{N}}$ are orthonormal, that is, $(\phi_i, \phi_j)_{L^2(\mathcal{D})} = \delta_{ij}$, and thus form a basis for $L^2(\mathcal{D})$; and (2) the corresponding eigenvalues form a sequence $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$, which in general decays with increasing mode number.

Because Γ is symmetric and positive semi-definite, it possesses the spectral decomposition [8, 41]

$$\Gamma(x, x') = \sum_{i=1}^{\infty} \lambda_i \phi_i(x) \phi_i(x'). \quad (10)$$

With these definitions, the KL expansion of a random field $\alpha \in L^2(\mathcal{D} \times \Theta)$ is defined by the following series [26]

$$\alpha(x, \theta) = \mu(x) + \sum_{i=1}^{\infty} \sqrt{\lambda_i} \phi_i(x) \xi_i(\theta), \quad \text{where} \quad \xi_i(\theta) := \frac{1}{\sqrt{\lambda_i}} \int_{\mathcal{D}} (\alpha(x, \theta) - \mu(x)) \phi_i(x) dx. \quad (11)$$

While $\{\phi_i\}_{i \in \mathbb{N}}$ are pairwise L^2 -orthogonal on \mathcal{D} , the $\{\xi_i\}_{i \in \mathbb{N}}$ are pairwise uncorrelated zero-mean random variables [14]. For this reason the KL expansion is sometimes said to be bi-orthogonal.

3.2. Truncated KL expansion

In order to represent a random field in a discrete numerical computation it is necessary to discretize the continuous probability space. This can be achieved by truncating the KL expansion after M terms and thus reducing the dimension of the stochastic space to M uncorrelated random variables

$$\tilde{\alpha}_M(x, \theta) = \mu(x) + \sum_{i=1}^M \sqrt{\lambda_i} \phi_i(x) \xi_i(\theta). \quad (12)$$

The mean of a random field is not affected by the discretization. The variance of the discretization on the other hand can be derived from the spectral decomposition in equation (10)

$$\mathbb{E}[(\tilde{\alpha}_M(x, \theta) - \mu(x))^2] = \sum_{i=1}^M \lambda_i \phi_i^2(x). \quad (13)$$

The variance of the discretized random field converges uniformly in \mathcal{D} and in $L^2(\Theta, \Sigma, \mathbb{P})$ towards the true variance [47]

$$\lim_{M \rightarrow \infty} \sum_{i=1}^M \lambda_i \phi_i^2(x) = \Gamma(x, x). \quad (14)$$

Furthermore, it can be shown that the KL expansion is optimal with respect to the global mean-squared error among all series expansions of truncation order M [14].

3.3. Variational formulation

The variational formulation or weak form of the integral eigenvalue problem introduced in equation (9) states

Find $\{\lambda, \phi\} \in \mathbb{R}_0^+ \times L^2(\mathcal{D})$ such that

$$\int_{\mathcal{D}} \left(\int_{\mathcal{D}'} \Gamma(x, x') \phi(x') \, dx' - \lambda \phi(x) \right) \psi(x) \, dx = 0 \quad \forall \psi \in L^2(\mathcal{D}). \quad (15)$$

Confining the solution to the finite-dimensional subspace $\mathcal{S}_h \subset L^2(\mathcal{D})$ we obtain the discrete variational formulation

Find $\{\lambda, \phi\} \in \mathbb{R}_0^+ \times \mathcal{S}_h$ such that

$$\int_{\mathcal{D}} \left(\int_{\mathcal{D}'} \Gamma(x, x') \phi_h(x') \, dx' - \lambda_h \phi_h(x) \right) \psi_h(x) \, dx = 0 \quad \forall \psi_h \in \mathcal{S}_h. \quad (16)$$

This is the Galerkin method for the homogeneous Fredholm integral eigenvalue problem of the second kind [3, 47]. Within the trial space under consideration, the Galerkin method produces the best L^2 fit of the analytical modes. The resulting discrete modes preserve exactly the L^2 orthogonality property of the analytical mode-shapes. Furthermore, it can be shown that a variational treatment using the Galerkin method leads to eigenvalues that converge monotonically, under mesh refinement, towards the true eigenvalues [14].

3.4. Choice of the trial space

The choice of the trial space \mathcal{S}_h provides some freedom in the design of the Galerkin method. The recently proposed isogeometric Galerkin method for the KL expansion of random fields uses NURBS for the test and trial spaces [47]. This choice is motivated by the fact that the geometrical mapping is defined using NURBS and it is natural to remain within the isoparametric paradigm. This method shares the same technical challenges as all classical Galerkin methods applied to this class of problems [1, 47]: the formation and assembly costs, which have a time complexity of $\mathcal{O}(N_e^2 \cdot p^{3d})$, as well as the storage requirements,

which have space complexity of $\mathcal{O}(N^2)$, become quickly intractable with increasing number of elements N_e , polynomial degree p , dimension d and number of degrees of freedom N . A practical Galerkin method must address these difficulties in the design of the method.

We abandon the isoparametric concept and choose a different space to represent the finite-dimensional solution. Our choice offers multiple computational advantages without sacrificing higher-order accuracy and robustness. We define the trial space for the Galerkin method as

$$\mathcal{S}_h := \text{span} \left\{ \frac{B_i(\hat{x})}{\sqrt{\det DF(\hat{x})}} \right\}_{i \in \mathcal{I}}. \quad (17)$$

Because the geometrical mapping F is smooth and invertible the Jacobian determinant is never singular, that is, $\det DF(\hat{x}) > 0$ for all $\hat{x} \in \hat{\mathcal{D}}$. Importantly, the functions are linearly independent due to linear independence of B-splines. In general, however, these basis functions will not form a partition of unity. Instead, the characterizing property is that products of these functions are integral preserving, that is, they transform as volume forms

$$\int_{\mathcal{D}} \frac{B_i(\hat{x})}{\sqrt{\det DF(\hat{x})}} \frac{B_j(\hat{x})}{\sqrt{\det DF(\hat{x})}} dx = \int_{\mathcal{D}} \frac{B_i(\hat{x})B_j(\hat{x})}{\det DF(\hat{x})} dx = \int_{\hat{\mathcal{D}}} B_i(\hat{x})B_j(\hat{x}) d\hat{x}.$$

3.5. Matrix formulation

After substituting the desired subspace for the test and trial functions and performing minor algebraic manipulations, the discretized Galerkin method results in a generalized algebraic eigenvalue problem

$$\mathbf{A}\mathbf{v}_h = \lambda_h \mathbf{Z}\mathbf{v}_h, \quad (18)$$

where the system matrices are formed by evaluating

$$\begin{aligned} A_{ij} &= \int_{\hat{\mathcal{D}}} \int_{\hat{\mathcal{D}'}} \Gamma(x(\hat{x}), x(\hat{x}')) \frac{B_i(\hat{x})}{\sqrt{\det DF(\hat{x})}} \frac{B_j(\hat{x}')}{\sqrt{\det DF(\hat{x}')}} \det DF(\hat{x}) \det DF(\hat{x}') d\hat{x}' d\hat{x} \\ &= \int_{\hat{\mathcal{D}}} \int_{\hat{\mathcal{D}'}} \Gamma(x(\hat{x}), x(\hat{x}')) B_i(\hat{x}) B_j(\hat{x}') \sqrt{\det DF(\hat{x}) \det DF(\hat{x}')} d\hat{x}' d\hat{x} \end{aligned} \quad (19)$$

and

$$\begin{aligned} Z_{ij} &= \int_{\hat{\mathcal{D}}} \frac{B_i(\hat{x})}{\sqrt{\det DF(\hat{x})}} \frac{B_j(\hat{x})}{\sqrt{\det DF(\hat{x})}} \det DF(\hat{x}) d\hat{x} \\ &= \int_{\hat{\mathcal{D}}} B_i(\hat{x}) B_j(\hat{x}) d\hat{x}. \end{aligned} \quad (20)$$

As a result of the chosen solution space, the mass matrix \mathbf{Z} has a Kronecker structure and can be decomposed into $k = 1, \dots, d$ univariate mass matrices

$$\mathbf{Z}_k := Z_{i_k j_k} = \int_0^1 B_{i_k, p_k}(\hat{x}_k) B_{j_k, p_k}(\hat{x}_k) d\hat{x}_k, \quad i_k, j_k = 1, \dots, n_k. \quad (21)$$

The system mass matrix \mathbf{Z} can be then written as

$$\mathbf{Z} = \mathbf{Z}_d \otimes \dots \otimes \mathbf{Z}_1. \quad (22)$$

Instead of computing and storing the matrix \mathbf{Z} , we precompute and store the matrices \mathbf{Z}_k , $k = 1, \dots, d$. Furthermore, it is the Kronecker structure that allows us to inexpensively reformulate the generalized eigenvalue problem to a standard algebraic eigenvalue problem.

Remark 3.2. In practice Z_k in (21) is computed exactly up to machine precision using Gauss-Legendre numerical quadrature with $p+1$ quadrature points per element, where p is the polynomial degree in component direction k . Because the domain of integration is one-dimensional the formation and assembly costs of $\mathcal{O}(n_e p^3)$ as well as the storage costs of $\mathcal{O}(pn)$ bytes are negligible compared to the total solver costs. Here n_e is the number of univariate elements and n is the univariate number of degrees of freedom in component direction k .

4. Efficient matrix-free solution strategy

There are two major challenges when applying the Galerkin method to discretize the homogeneous Fredholm integral eigenvalue problem in (9). Firstly, the variational formulation requires integration over a $2d$ -dimensional domain to evaluate the matrix entries in \mathbf{A} . This leads to formation and assembly costs with complexity $\mathcal{O}(N_e^2 p^{3d})$, where N_e is the global number of elements, p the polynomial degree and d the spatial dimension. Secondly, because the matrix is dense, \mathbf{A} requires insurmountable memory storage for any practical problem of interest. Several techniques have been presented in the literature in order to deal with these challenges, for example by approximation with low-rank matrices like the hierarchical matrices [1, 16, 29] or by using Fast Multipole Methods [52]. In this work we present a combination of four techniques to deal with the aforementioned challenges:

1. Reformulation of the generalized eigenvalue problem into an equivalent standard eigenvalue problem;
2. Interpolation based quadrature for variational formulations of integral equations;
3. Efficient formation of finite element arrays based on Kronecker matrix-vector product;
4. Formulation of a matrix-free and multi-threaded matrix-vector product for the Lanczos algorithm.

The reformulation into a standard algebraic eigenvalue problem significantly reduces the computational cost and simultaneously improves conditioning. By exploiting the Kronecker structure of the right-hand-side mass matrix we can perform this reformulation with negligible overhead. The proposed non-standard quadrature technique that we call *interpolation based integration* is tailored for variational formulations of integral equations. The technique is optimal in the sense that few quadrature points are required while integrating a rich space of tensor product functions on the $2d$ -dimensional domain $\mathcal{D} \times \mathcal{D}$. Importantly, the technique lends itself to multidimensional tensor contraction due to the Kronecker structure of the involved matrices. This significantly speeds up the evaluation of integrals over high-dimensional domains and scales favorably with polynomial degree. Finally, all techniques are combined within a matrix-free evaluation scheme that is embarrassingly parallel and requires minimal memory storage. The formation and assembly costs of our approach are negligible compared to the remaining solver costs of the Lanczos eigenvalue solver, which is $\mathcal{O}(\tilde{N}^2 \cdot N_{\text{iter}}/N_{\text{thread}})$. Here \tilde{N} is the global number of degrees of freedom of the interpolation space, N_{iter} is the number of iterations of the eigensolver and N_{thread} is the number of simultaneous processes. In the following we discuss each of the proposed techniques in more detail.

4.1. Reformulation into a standard algebraic eigenvalue problem

Let us consider a Cholesky factorization of the mass matrix $\mathbf{Z} = \mathbf{L}\mathbf{L}^\top$ and define a linear transformation of the eigenvectors $\mathbf{v}'_h := \mathbf{L}^\top \mathbf{v}_h$. The generalized eigenvalue problem can then be rewritten (see [49, Chapter 9.2.2]) as a standard eigenvalue problem with unchanged eigenvalues corresponding to new eigenvectors \mathbf{v}'_h

$$\mathbf{A}' \mathbf{v}'_h = \lambda_h \mathbf{v}'_h, \quad \text{where } \mathbf{A}' := \mathbf{L}^{-1} \mathbf{A} \mathbf{L}^{-\top} \quad (23)$$

It is expected that the new system matrix \mathbf{A}' has improved conditioning compared to \mathbf{A} . The kernel is positive-definite and symmetric. In practice, it often quickly tends to zero for increasing distance $\|x - x'\|_2$. In the limiting case, where $\Gamma(x, x') \rightarrow \delta(x, x')$, the system matrix $\mathbf{A} \rightarrow \mathbf{Z}$ and hence the preconditioner would be ideal.

Although improved conditioning is beneficial, the main reason for the chosen transformation is efficiency. Solution of a standard algebraic eigenvalue problem is much less expensive than solution of a generalized

eigenvalue problem. The transformation itself is inexpensive. Using the Kronecker structure of Z and the properties (2b) and (2d) we may write

$$\begin{aligned} Z &= Z_d \otimes \dots \otimes Z_1 \\ &= L_d L_d^\top \otimes \dots \otimes L_1 L_1^\top \\ &= (L_d \otimes \dots \otimes L_1)(L_d \otimes \dots \otimes L_1)^\top = LL^\top. \end{aligned} \quad (24)$$

Here $L_k L_k^\top$, $k = 1, \dots, d$, denote the Cholesky factorizations corresponding to the univariate mass matrices Z_k . Hence, instead of performing the Cholesky factorization for the complete system matrix $Z \in \mathbb{R}^{N \times N}$, which is the standard procedure in most solvers for generalized algebraic eigenvalue problems, we merely need the Cholesky factorizations for $Z_k \in \mathbb{R}^{n_k \times n_k}$, $k = 1, \dots, d$.

The factorization is precomputed once before using it in the eigenvalue solver. The associated computational cost is reduced from $\mathcal{O}(N^3)$ to $\mathcal{O}(n^3)$ flops, where $n = \max(n_1, \dots, n_d)$ and $N = n_1 \cdot \dots \cdot n_d$. Subsequently, the cost of applying the factorization in a single iteration of the eigenvalue solver is reduced from $\mathcal{O}(N^2)$ to $\mathcal{O}(n \cdot N)$. Besides a reduction in computational cost, this approach significantly reduces the required memory storage.

4.2. Interpolation based quadrature for integral equations

One of the most straightforward ways of improving efficiency is to design quadrature rules that require fewer evaluation points. This is especially true for the variational formulation of the Fredholm integral equation which requires numerical integration over a $2d$ -dimensional domain. In practice, accurate and efficient quadrature rules are designed as follows. First, one chooses a space of functions \mathbb{T} , called the target space for numerical quadrature, whose elements should be exactly integrated by the new quadrature rule. If this space is in some sense rich enough then the error due to quadrature can be bounded by the discretization error, which is needed to show optimal rates of convergence of the numerical method (see [18]). The next objective is to find a quadrature rule that requires as few points as possible to integrate all functions in \mathbb{T} .

We present a non-standard quadrature technique that generalizes the approach to quadrature presented for linear finite elements in [28, 29] to higher order splines. Our approach is tailored toward evaluating integrals found in variational formulations of integral equations and achieves a very low number of evaluation points while integrating exactly a rich space of functions. The target space \mathbb{T} is chosen such that the quadrature scheme exactly evaluates the integral

$$\int_{\hat{\mathcal{D}}} \int_{\hat{\mathcal{D}}'} \tilde{G}(\hat{x}, \hat{x}') B_i(\hat{x}) B_j(\hat{x}') d\hat{x}' d\hat{x}, \quad \tilde{G} \in \tilde{\mathcal{B}}_h(\hat{\mathcal{D}}) \otimes \tilde{\mathcal{B}}_h(\hat{\mathcal{D}}'). \quad (25)$$

using \tilde{N}^2 points. Here $\tilde{\mathcal{B}}_h$ is another d -dimensional spline space that can be chosen independently of \mathcal{B}_h and \tilde{N} is its dimension. In practice this space can be chosen to fit well with the integrand in the variational formulation of the integral equation. This provides additional flexibility to the quadrature scheme.

Because $\tilde{G} \in \tilde{\mathcal{B}}_h(\hat{\mathcal{D}}) \otimes \tilde{\mathcal{B}}_h(\hat{\mathcal{D}}')$ is a real-valued $2d$ -variate spline function it can be expanded in terms of B-spline basis functions and real-valued coefficients $\{\tilde{G}_{kl}\}_{k,l \in \tilde{\mathcal{I}}}$ as

$$\tilde{G}(\hat{x}, \hat{x}') := \sum_{k,l \in \tilde{\mathcal{I}}} \tilde{G}_{kl} \tilde{B}_k(\hat{x}) \tilde{B}_l(\hat{x}') \quad \text{with} \quad k, l \in \tilde{\mathcal{I}} := \{(i_1, \dots, i_d) : 1 \leq i_k \leq \tilde{n}_k\}.$$

Comparing the multidimensional integrand in (19) with the one in (25) we may conclude that the degrees of freedom $\{\tilde{G}_{kl}\}_{k,l \in \tilde{\mathcal{I}}}$ should be chosen such that \tilde{G} is a good approximation of the function $G : \hat{\mathcal{D}} \times \hat{\mathcal{D}} \mapsto \mathbb{R}^+$ defined as

$$G(\hat{x}, \hat{x}') := \hat{\Gamma}(\hat{x}, \hat{x}') \sqrt{\det DF(\hat{x}) \det DF(\hat{x}')} \quad (26)$$

Here $\hat{\Gamma}(\hat{x}, \hat{x}')$ is the pull-back of the kernel $\Gamma(x, x')$ from the physical to the parametric space using the geometrical mapping F .

The approximation \tilde{G} can be estimated in different ways. We follow a similar approach to the degenerate kernel approximation in [2] and choose to collocate G at the Greville abscissa [10]. This approach is both simple and combines high order accuracy with a minimal number of evaluation points. Let $\tilde{\mathbf{B}} = \tilde{B}_{ij} := \tilde{B}_j(\hat{x}_i)$ denote the d -variate spline collocation matrix evaluated at the Greville abscissa $\hat{x}_i \in \hat{\mathcal{D}}$, $i \in \tilde{\mathcal{I}}$. The interpolation problem states

Find $\{\tilde{G}_{kl}\}_{k,l \in \tilde{\mathcal{I}}}$ such that

$$\sum_{k,l \in \tilde{\mathcal{I}}} \tilde{G}_{kl} \tilde{B}_k(\hat{x}_i) \tilde{B}_l(\hat{x}'_j) = G(\hat{x}_i, \hat{x}'_j) \quad \forall i, j \in \tilde{\mathcal{I}} \quad (27)$$

This is equivalent to the matrix problem $\mathbf{G} = \tilde{\mathbf{B}} \tilde{\mathbf{G}} \tilde{\mathbf{B}}^\top$. Hence, the matrix of coefficients can be computed as $\tilde{\mathbf{G}} = \tilde{\mathbf{B}}^{-1} \mathbf{G} \tilde{\mathbf{B}}^{-\top}$. The computational cost of the interpolation can be significantly reduced from $\mathcal{O}(\tilde{N}^3)$ to $\mathcal{O}(\tilde{n} \cdot \tilde{N})$ flops, where $\tilde{n} = \max(\tilde{n}_1, \dots, \tilde{n}_d)$, by exploiting the Kronecker structure of $\tilde{\mathbf{B}}$. We decompose $\tilde{\mathbf{B}}$ into d univariate collocation matrices $\tilde{\mathbf{B}}_k$, $k = 1, \dots, d$, and use property (2c) to write its inverse as

$$\tilde{\mathbf{B}}^{-1} = \tilde{\mathbf{B}}_d^{-1} \otimes \dots \otimes \tilde{\mathbf{B}}_1^{-1} \quad \text{with} \quad \tilde{\mathbf{B}}_k := \tilde{B}_{i_k j_k} = \tilde{B}_{j_k, \tilde{p}_k}(\hat{x}_{i_k}) \quad (28)$$

In practice we compute d LU factorizations, each corresponding to a univariate matrix $\tilde{\mathbf{B}}_k$, to apply the inverse of $\tilde{\mathbf{B}}$ to a vector. Note, that this approach is similar to the approach we took in equation (24) for the Cholesky factorization of \mathbf{Z} .

4.3. Matrix formation

By substituting G in equation (19) with \tilde{G} we can approximate matrix \mathbf{A} by a matrix $\tilde{\mathbf{A}}$ with entries

$$\begin{aligned} \tilde{A}_{ij} &:= \int_{\hat{\mathcal{D}}} \int_{\hat{\mathcal{D}}'} \tilde{G}(\hat{x}, \hat{x}') B_i(\hat{x}) B_j(\hat{x}') d\hat{x}' d\hat{x} \\ &= \sum_{k,l \in \tilde{\mathcal{I}}} \tilde{G}_{kl} \int_{\hat{\mathcal{D}}} \int_{\hat{\mathcal{D}}'} \tilde{B}_k(\hat{x}) \tilde{B}_l(\hat{x}') B_i(\hat{x}) B_j(\hat{x}') d\hat{x}' d\hat{x} \\ &= \sum_{k,l \in \tilde{\mathcal{I}}} \tilde{G}_{kl} \int_{\hat{\mathcal{D}}} \tilde{B}_k(\hat{x}) B_i(\hat{x}) d\hat{x} \int_{\hat{\mathcal{D}}'} \tilde{B}_l(\hat{x}') B_j(\hat{x}') d\hat{x}' \end{aligned}$$

Hence, using the tensor product structure of the interpolation space we have separated the $2d$ -dimensional integral into a product of two d -dimensional integrals. In matrix notation we may write $\tilde{\mathbf{A}} = \mathbf{M}^\top \tilde{\mathbf{G}} \mathbf{M}$, or

$$\tilde{A}_{ij} = \sum_{k,l \in \tilde{\mathcal{I}}} \tilde{G}_{kl} M_{ki} M_{lj} \quad (29)$$

Here $\mathbf{M} := M_{ij}$ is a mass matrix

$$M_{ij} = \int_{\hat{\mathcal{D}}} \tilde{B}_i(\hat{x}) B_j(\hat{x}) d\hat{x}. \quad (30)$$

Similarly, as we did for matrix \mathbf{Z} in (22), we can exploit the Kronecker structure and decompose \mathbf{M} into d univariate mass matrices $\mathbf{M}_k := M_{i_k j_k}$

$$\mathbf{M} = \mathbf{M}_1 \otimes \dots \otimes \mathbf{M}_d \quad \text{with} \quad M_{i_k j_k} = \int_0^1 \tilde{B}_{i_k, \tilde{p}_k}(\hat{x}_k) B_{j_k, p_k}(\hat{x}_k) d\hat{x}_k \quad (31)$$

As in the case of \mathbf{Z}_k , $k = 1, \dots, d$, these univariate matrices are computed up to machine precision as discussed in remark 3.2. The approximation error $\mathbf{A} - \tilde{\mathbf{A}}$ is entirely due to the interpolation error $G - \tilde{G}$. Hence, accurate approximation of G should result in an accurate approximation of \mathbf{A} .

4.4. Matrix-free solution strategy

The interpolation based quadrature technique introduced in the previous section involves computation of the matrix of coefficients $\tilde{\mathbf{G}} := \tilde{G}_{kl}$. This matrix is dense and has N^2 entries. Consequently, storage of $\tilde{\mathbf{G}}$ is just as inconvenient as storing $\tilde{\mathbf{A}}$ and becomes quickly intractable with problem size. In this section we propose a matrix-free evaluation of the matrix-vector product $\mathbf{v}' \mapsto \tilde{\mathbf{A}}\mathbf{v}'$ that does not require explicit access to matrix $\tilde{\mathbf{G}}$ or $\tilde{\mathbf{A}}$.

4.4.1. Basic setup

We have the following standard algebraic eigenvalue problem

$$\tilde{\mathbf{A}}'\mathbf{v}' = \lambda_h\mathbf{v}'. \quad (32a)$$

Here, the system matrix $\tilde{\mathbf{A}}'$ can be written as

$$\tilde{\mathbf{A}}' = \mathbf{L}^{-1}\mathbf{M}^\top\tilde{\mathbf{B}}^{-1}\mathbf{J}\Gamma\mathbf{J}\tilde{\mathbf{B}}^{-\top}\mathbf{M}\mathbf{L}^{-\top}, \quad (32b)$$

where \mathbf{J} is a diagonal matrix with diagonal entries given by the square roots of Jacobian determinants evaluated at the Greville abscissa, and \mathbf{B} , \mathbf{M} and \mathbf{L} are all Kronecker product matrices. Consequently, a matrix-vector product with any of these matrices can be performed close to linear time complexity. The matrix vector product $\mathbf{v}' \mapsto \tilde{\mathbf{A}}'\mathbf{v}'$ can be subdivided into the following operations

$$\Gamma := \hat{\Gamma}(\hat{x}_k, \hat{x}'_l) \quad (\text{Evaluation of the kernel at the Greville abscissa}) \quad (32c)$$

$$\mathbf{G} = \mathbf{J}\Gamma\mathbf{J} \quad (\text{Scaling of the kernel}) \quad (32d)$$

$$\tilde{\mathbf{G}} = \tilde{\mathbf{B}}^{-1}\mathbf{G}\tilde{\mathbf{B}}^{-\top} \quad (\text{Interpolation of the scaled kernel}) \quad (32e)$$

$$\tilde{\mathbf{A}} = \mathbf{M}^\top\tilde{\mathbf{G}}\mathbf{M} \quad (\text{Evaluation of the integrals}) \quad (32f)$$

$$\tilde{\mathbf{A}}' = \mathbf{L}^{-1}\tilde{\mathbf{A}}\mathbf{L}^{-\top} \quad (\text{Application of the preconditioner}) \quad (32g)$$

In the following subsection we present a matrix-free matrix-vector product that incorporates each of the above steps. Except for the diagonal matrix \mathbf{J} , none of the above matrices are stored explicitly. Only the corresponding univariate matrices are stored and used in the Kronecker products, while the entries of Γ are computed on the fly.

4.4.2. Matrix-free algorithm

Let $\hat{\Gamma}_{k_1\dots k_d l_1\dots l_d} := \hat{\Gamma}(\hat{x}_{1,k_1}, \dots, \hat{x}_{d,k_d}, \hat{x}'_{1,l_1}, \dots, \hat{x}'_{d,l_d}) \in \mathbb{R}^{\tilde{n}_1 \times \dots \times \tilde{n}_d \times \tilde{n}_1 \times \dots \times \tilde{n}_d}$ denote the function values of the kernel evaluated at the tensor product grid of the Greville abscissa in $\hat{\mathcal{D}} \times \hat{\mathcal{D}}$. The proposed evaluation order of the matrix-free matrix-vector product is summarized in Algorithm 1.

Algorithm 1 Matrix-free evaluation of the matrix-vector product $\mathbf{v}' \mapsto \tilde{\mathbf{A}}'\mathbf{v}'$

Input: $v_{i_1\dots i_d} \in \mathbb{R}^{n_1 \times \dots \times n_d}$, $J_{l_1\dots l_d} \in \mathbb{R}^{\tilde{n}_1 \times \dots \times \tilde{n}_d}$, $B_{i_k j_k} \in \mathbb{R}^{\tilde{n}_k \times \tilde{n}_k}$ and $M_{l_k j_k} \in \mathbb{R}^{\tilde{n}_k \times n_k}$

Output: $v'_{i_1\dots i_d} \in \mathbb{R}^{n_1 \times \dots \times n_d}$

- 1: $V_{j_1\dots j_d} \leftarrow L_{i_1 j_1}^{-1} \dots L_{i_d j_d}^{-1} v_{i_1\dots i_d}$ ▷ Preconditioning from right
 - 2: $X_{k_1\dots k_d} \leftarrow M_{k_1 j_1} \dots M_{k_d j_d} V_{j_1\dots j_d}$
 - 3: $Y_{l_1\dots l_d} \leftarrow B_{k_1 l_1}^{-1} \dots B_{k_d l_d}^{-1} X_{k_1\dots k_d}$ ▷ Use LU-factorization of B_k , $k = 1, \dots, d$
 - 4: $Y'_{l_1\dots l_d} \leftarrow J_{l_1\dots l_d} \odot Y_{l_1\dots l_d}$
 - 5: $Z'_{k_1\dots k_d} \leftarrow \hat{\Gamma}_{k_1\dots k_d l_1\dots l_d} Y'_{l_1\dots l_d}$ ▷ Evaluate in parallel without forming Γ
 - 6: $Z_{k_1\dots k_d} \leftarrow J_{k_1\dots k_d} \odot Z'_{k_1\dots k_d}$
 - 7: $Y_{j_1\dots j_d} \leftarrow B_{j_1 k_1}^{-1} \dots B_{j_d k_d}^{-1} Z_{k_1\dots k_d}$ ▷ Use LU-factorization of B_k , $k = 1, \dots, d$
 - 8: $V_{l_1\dots l_d} \leftarrow M_{j_1 l_1} \dots M_{j_d l_d} Y_{j_1\dots j_d}$
 - 9: $v'_{i_1\dots i_d} \leftarrow L_{i_1 l_1}^{-1} \dots L_{i_d l_d}^{-1} V_{l_1\dots l_d}$ ▷ Preconditioning from left
-

The matrix-free matrix vector product $\mathbf{v}' \mapsto \tilde{\mathbf{A}}'\mathbf{v}'$ is evaluated in nine separate stages. Stage one applies back-substitution of the upper triangular matrix \mathbf{L}^\top and exploits its Kronecker structure. Stage two applies a matrix vector product with matrix \mathbf{M} and again exploits its Kronecker structure. Stage three applies back-substitution using the factorization of the interpolation matrix \mathbf{B} . Again, Kronecker structure is essential to reduce both the space and time complexity of the back-substitution. In stage four the coefficient vector is element-wise multiplied by the square root of the Jacobian determinant evaluated at the Greville abscissa. Here, element-wise multiplication is denoted by the \odot symbol. Stage five dominates the computational cost of Algorithm 1. This stage represents a dense matrix-vector product. To perform this step without explicitly forming matrix Γ we compute its entries on the fly, one row at a time. We compute products of the coefficient vector with several rows of Γ in parallel. Stages six to nine are equivalent to stages four to one, due to the symmetry of the operator.

Due to the iterative solution process, the matrix-vector product needs to be evaluated at each iteration. The number of iterations is dependent on the number of required eigenmodes, the conditioning of the algebraic eigenvalue problem and the efficiency of the eigensolver. In this work we use the standard implicitly restarted Lanczos method [15].

5. Computational complexity analysis

The goal of a time-complexity analysis is to obtain an estimate of the computational cost that scales linearly with time. This cost is expressed in terms of certain parameters that depend on the problem size, the dimension and the polynomial degree. For this purpose, let us introduce the following notation:

| | |
|----------------------------|---|
| n | number of degrees of freedom of the trial space in one component direction; |
| \tilde{n} | number of degrees of freedom of the interpolation space in one component direction; |
| $N := n^d$ | total number of degrees of freedom of the trial space; |
| $\tilde{N} := \tilde{n}^d$ | total number of degrees of freedom of the interpolation space; |
| N_e | total number of spatial elements in the trial space; |
| N_q | number of quadrature points in a standard quadrature loop. |
| N_{iter} | number of iterations of the matrix-free algorithm; |
| N_{thread} | number of simultaneous shared memory processes in the matrix-vector product. |

5.1. Standard finite element procedures

In the following we present the computational complexity of standard finite element procedures for higher-order finite elements. We use the tensor product structure of the high-dimensional space $\hat{\mathcal{D}} \times \hat{\mathcal{D}}$ to minimize the involved computations. The general setting for this analysis is (1) $\hat{\mathcal{D}} \times \hat{\mathcal{D}}$ has N_e^2 elements; and (2) we assume a quadrature rule $Q(f) := \sum_{k=1}^{N_q} w_k f(x_k)$, with $1 \leq N_q \leq (p+1)^d$, to integrate the products on every d -dimensional element \square^d in $\hat{\mathcal{D}}$.

The leading term in formation and assembly is determined by the cost of forming the element matrices. Consider the following element matrix

$$\begin{aligned}
A_{ij}^e &= \int_{\square^d} B_i(\hat{x}) \int_{\square^d} \Gamma(\hat{x}, \hat{x}') B_j(\hat{x}') d\hat{x}' d\hat{x} \\
&\approx \sum_{k=1}^{N_q} w_k B_i(\hat{x}_k) \sum_{l=1}^{N_q} w_l \Gamma(\hat{x}_k, \hat{x}'_l) B_j(\hat{x}'_l) \\
&= \sum_{k=1}^{N_q} C_{ik} D_{kj} \quad \text{with} \quad D_{kj} = \sum_{l=1}^{N_q} w_l \Gamma(\hat{x}_k, \hat{x}'_l) B_j(\hat{x}'_l)
\end{aligned}$$

with $i, j \in \mathcal{I}$. We see that A_{ij}^e can be formed by the matrix product of matrices $\mathbf{C} \in \mathbb{R}^{(p+1)^d \times N_q}$ and $\mathbf{D} \in \mathbb{R}^{N_q \times (p+1)^d}$. This matrix product costs $\mathcal{O}(N_q(p+1)^{2d})$. The formation of C_{ik} is negligible. The formation

of D_{kj} on the other hand is $\mathcal{O}(N_q^2(p+1)^d)$. Since $N_q \leq (p+1)^d$ the leading term is $N_q(p+1)^{2d}$. Hence, the total cost of forming one element matrix is $\mathcal{O}(N_q(p+1)^{2d})$. In total we have to integrate over all N_e^2 multidimensional elements of $\hat{\mathcal{D}} \times \hat{\mathcal{D}}$. With that, the total cost of forming \mathbf{A} is $\mathcal{O}(N_e^2 N_q(p+1)^{2d})$. Using a Gauss-Legendre quadrature rule with $(p+1)$ quadrature points in every coordinate direction gives in total $N_q = (p+1)^d$ quadrature points, and we can expect a leading cost proportional to $\mathcal{O}(N_e^2(p+1)^{3d})$.

5.2. Finite element procedures employing sum factorization

Estimates presented in the previous subsection hold for classical hp finite element procedures employing a standard quadrature loop. Next we discuss the complexity of finite element procedures that employ sum factorization instead of a standard quadrature loop. Sum factorization significantly speeds up the element array formation by exploiting the tensorial structure of both the finite element basis and the used quadrature rules [42, 55, 57]. It's worth noting that, due to the structure of the integral operator, the sum factorization technique looks somewhat different than is standard in the hp finite element method.

The setting for this analysis is (1) $\hat{\mathcal{D}} \times \hat{\mathcal{D}}$ has N_e^2 rectangular elements; (2) we use a tensor product basis of polynomial degree p on every element; and (3) we use a tensor product of univariate Gauss-Legendre quadrature rules $Q(f) := \sum_{k=1}^{p+1} \omega_k f(x_k)$ to integrate the products on every d -dimensional element \square^d in $\hat{\mathcal{D}}$. Consider the element matrix

$$\begin{aligned} A_{ij}^e &= \int_{\square^d} B_i(x) \int_{\square^d} \Gamma(x, x') B_j(x') dx' dx \\ &\approx \sum_{k_1=1}^{p+1} B_{i_1,p}(x_{1,k_1}) \sum_{k_2=1}^{p+1} B_{i_2,p}(x_{2,k_2}) \cdots \sum_{k_d=1}^{p+1} B_{i_d,p}(x_{d,k_d}) \\ &\quad \sum_{l_1=1}^{p+1} B_{j_1,p}(x'_{1,l_1}) \sum_{l_2=1}^{p+1} B_{j_2,p}(x'_{2,l_2}) \cdots \sum_{l_d=1}^{p+1} \Gamma(x_{1,k_1}, \dots, x_{d,k_d}, x'_{1,l_1}, \dots, x'_{d,l_d}) B_{j_d,p}(x'_{d,l_d}) \end{aligned}$$

with $i, j \in \mathcal{I}$. Sum factorization is essentially tensor contraction. The kernel evaluated at the grid of quadrature points is a tensor $\Gamma_{k_1 \dots k_d l_1 \dots l_d} \in \mathbb{R}^{(p+1) \times \dots \times (p+1)}$ and is contracted with matrices $B_{j_z,p}(x'_{z,l_z}) \in \mathbb{R}^{(p+1) \times (p+1)}$, for $z = 1, \dots, d$, and subsequently with matrices $B_{i_z,p}(x_{z,k_z}) \in \mathbb{R}^{(p+1) \times (p+1)}$ for $z = 1, \dots, d$. The cost of every contraction is $\mathcal{O}((p+1)^{2d+1})$ flops. In total there are $2d$ such tensor contractions. Hence, the element matrix formation cost for \mathbf{A}^e is $\mathcal{O}(2d(p+1)^{2d+1})$ flops. With N_e^2 elements, the leading cost of forming \mathbf{A} is $\mathcal{O}(2dN_e^2(p+1)^{2d+1})$ flops.

5.3. Proposed strategy using interpolation based quadrature

In order to analyze the computational complexity of the proposed solution strategy we must address each stage of the matrix-free matrix-vector product introduced in Algorithm 1. Let us consider the complexity in one iteration of the matrix-free algorithm.

- Stage 1 has a cost $\mathcal{O}(dn^{d+1})$.
- Stage 2 has a cost depending on the chosen projection space,
 - for $n > \tilde{n}$ the cost is $\mathcal{O}(dpn^d)$,
 - for $n = \tilde{n}$ the cost is $\mathcal{O}(dpn^d)$,
 - for $n < \tilde{n}$ the cost is $\mathcal{O}(dp\tilde{n}^d)$.
- Stage 3 has a cost $\mathcal{O}(d\tilde{n}^{d+1})$.
- Stage 4 has a cost $\mathcal{O}(\tilde{n}^d)$.
- Stage 5 has a cost $\mathcal{O}(\tilde{N}^2/N_{\text{thread}})$.

The remaining steps 6, 7, 8 and 9 are equivalent to steps 4, 3, 2 and 1. In step 2 we assume sparse matrix algebra. In dense algebra, p can be replaced by n . In steps 1 and 3 we assume that the Cholesky and LU

factorizations of the univariate matrices of size $n \times n$ and $\tilde{n} \times \tilde{n}$, respectively, have been precomputed and are available. Subsequently, the solver costs are $\mathcal{O}(n^2)$ and $\mathcal{O}(\tilde{n}^2)$ flops, respectively, for each application of the factorization. The cost of a single iteration is typically dominated by step 5, which does not depend on the polynomial degree p . Fortunately, this step is embarrassingly parallel. Hence, the time complexity of the matrix free algorithm is $\mathcal{O}(\tilde{N}^2 N_{\text{iter}}/N_{\text{thread}})$ flops.

5.4. Storage comparison

Both matrix-free and non-matrix-free methods need to store the resulting eigenmodes. Storage of the results takes roughly $L \cdot 8N$ bytes, where L is the number of eigenmodes that need to be computed. Additionally, the standard approach that stores the dense left-hand-side system matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ requires storage of $N \times N$ floating point numbers. Using double precision floating point arithmetic the storage requirements are $8N^2$ bytes. The additional storage of the matrix-free approach scales linearly with problem size, with an asymptotic leading term of roughly $2 \cdot 8N$ bytes, again using double precision floating point arithmetic. If step 4 of the matrix-free algorithm is performed using shared memory parallelism then one can expect this to increase to $(1 + N_{\text{thread}}) \cdot 8N$ where N_{thread} is the number of simultaneous processes. Consequently, the storage cost of the matrix-free approach is typically governed by storage of the results and is thus optimal. Table 2 and 3 summarize the leading terms in storage of the two alternatives.

Table 2: The leading terms in storage costs of a method that explicitly stores the main system matrix. Here L refers to the number of eigenmodes that need to be stored. The storage cost is dominated by storage of the system matrix.

| Number of degrees of freedom | 10^3 | 10^4 | 10^5 | 10^6 |
|------------------------------|----------------|-----------------|------------------|----------------|
| Results storage | $L \cdot 8$ KB | $L \cdot 80$ KB | $L \cdot 800$ KB | $L \cdot 8$ MB |
| Matrix storage | 8 MB | 800 MB | 80 GB | 8 TB |

Table 3: The leading terms in storage costs of a matrix-free approach. Here L refers to the number of eigenmodes that need to be stored. The storage cost of the matrix-free method is typically dominated by the storage of the results.

| Number of degrees of freedom | 10^3 | 10^4 | 10^5 | 10^6 |
|------------------------------|----------------|-----------------|------------------|----------------|
| Results storage | $L \cdot 8$ KB | $L \cdot 80$ KB | $L \cdot 800$ KB | $L \cdot 8$ MB |
| Matrix-free approach | $2 \cdot 8$ KB | $2 \cdot 80$ KB | $2 \cdot 800$ KB | $2 \cdot 8$ MB |

6. Numerical results

In this section we present numerical results that illustrate the accuracy, robustness and computational efficiency of the proposed matrix-free Galerkin method. An accuracy study is performed in the case of one dimension. This study case gives insight into the accuracy attained by interpolation based quadrature of the covariance function and its affect on approximating the eigenvalue spectra. Several three-dimensional benchmarks illustrate the computational performance attained for a range of polynomial degrees and different refinement strategies of the interpolation as well as the solutions space. The computations are performed on two Intel(R) Xeon(R) CPU E5-2630 @ 2.30GHz (24 threads) and 12x8 GB of ECC DDR3L 1600MHz RAM. The Python implementation in this work relies heavily on the packages Numpy [56] and Scipy [53] for the linear algebra and solver functionalities. In order to achieve high performance, crucial parts of the code base are just-in-time compiled using the LLVM-based Python-compiler Numba [32, 33].

6.1. One-dimensional case study

Consider a one-dimensional random field defined on domain $\mathcal{D} = [0, 1] \subset \mathbb{R}$. We investigate

- (i) the relative $L^2(\mathcal{D})$ interpolation error of the kernel, $G - \tilde{G}$, with respect to uniform h -refinement enforcing C^{p-1} continuity across element boundaries. We consider the cases where G is the Gaussian and exponential kernel with $b = 0.1$, a characteristic domain length $L = 1$ and a variance $\sigma^2 = 1$.

- (ii) the normalized spectra corresponding to an exponential kernel with $b = 1$.

Remark 6.1. *Normalized spectra corresponding a Gaussian kernel cannot be reliably computed across the full range of eigenvalues because the smallest eigenvalues quickly approach zero up to machine precision.*

Kernel approximation

Figure 3a shows the convergence towards the Gaussian kernel for polynomial degrees 1 through 8. It seems that the even degrees $p = \{2, 4, 6, 8\}$ perform relatively better than the preceding odd degrees $p = \{1, 3, 5, 7\}$. It is evident that higher-order interpolation of a smooth kernel leads to a higher-order convergence rate in the approximation. The smooth interpolation space is not as suitable for approximation of kernels that have low regularity. Approximation of the exponential kernel in Figure 3b, which is C^0 along $x = x'$, shows that higher-order continuity of the basis does not lead to an increased convergence rate. This behavior is in agreement with convergence estimates for spline approximation of arbitrary smoothness [50].

Spectral approximation

Although the proposed method, in its current form, is best suited for smooth kernels like the Gaussian kernel, excellent approximation of the eigenvalues corresponding to non-smooth kernels is still possible. Figure 4 depicts the full spectrum corresponding to the exponential kernel with $b = 1$. The proposed Galerkin method using interpolation based quadrature (IBQ) is compared to the isogeometric Galerkin method proposed in [47] and a classical C^0 finite element solution in the case of polynomial degree $p = 2$. The interpolation space is set to $\tilde{h} = 0.005 \cdot bL$. The proposed method (IBQ) exhibits the same advantageous characteristics as the standard isogeometric Galerkin method [47] and exhibits no branching phenomena as in the case of the C^0 -continuous finite element approximation. Due to their increased continuity across element boundaries, splines achieve a higher accuracy per degree of freedom and increased robustness as compared with classical C^0 finite element methods. These results are in agreement with several other studies that have investigated spectral approximations corresponding to eigenvalue problems in structural mechanics [6, 22, 23, 46].


Remark 6.2. *The reference solution is computed using a standard isogeometric Galerkin method [47] with fifty thousand degrees of freedom. The first twenty eigenvalues have been validated up to machine accuracy by the analytical approach described in [14, Ch. 2.3.3, page 28-35]. The analytical computation of these eigenvalues involves solving for roots of a complex equation and is for that reason avoided beyond the first twenty eigenvalues.*

6.2. Random field with exponential kernel in a three-dimensional half-open cylindrical domain


In the first three-dimensional example we investigate a random field defined in a half-open cylindrical domain as shown in figure 5. The kernel is set to the exponential covariance function with a correlation length bL equal to the half of the characteristic length L . The variance of the random field is $\sigma^2 = 1$. We note that this example is also studied in [47].

Since the kernel is C^0 along $x = x'$, there is no gain in using a high-order interpolation space. Therefore, the chosen refinement strategy is simple uniform h -refinement in each parametric direction. The polynomial order in each direction is the minimal order, which allows an exact representation of the geometry. This implies quadratic splines in the circumferential direction and linear splines in the other directions. The solution space and the interpolation space meshes for each refined case are shown in Figure 6. The first twenty largest eigenvalues together with information about the mesh and computational cost are tabulated in Table 4. The first nine eigenfunctions corresponding to the nine largest eigenvalues are visualized in Figure 7 by weighting each eigenfunction by the square root of the corresponding eigenvalue.

The presolution formation and assembly time includes the formation and assembly of the univariate mass matrices, the interpolation matrices and their factorizations as well as the computation of the Jacobian determinants at the Greville abscissa. The results in Table 4 indicate that these setup costs are negligibly small compared to the total solution time of the Lanczos eigenvalue solver, which is dominated by the matrix-free evaluation of the matrix-vector product (step 5 in Algorithm 1). The number of iterations of



(a) Gaussian kernel



(b) Exponential kernel

Figure 3: Normalized L^2 interpolation error in the one-dimensional study case for multiple polynomial degrees. The error is given with respect to the mesh size of the interpolation space (bottom axis), as well as the mesh size of the interpolation space normalized by the correlation length (top axis). The convergence rates are approximately $\mathcal{O}(\tilde{h}^{p+1})$ in (a) and $\mathcal{O}(\tilde{h}^{3/2})$ in (b).




Figure 4: Ratio of the approximated eigenvalues to the reference eigenvalues over a full spectrum of 501 eigenmodes in the one-dimensional study case with an exponential kernel and a correlation length equal to the domain length.

the Lanczos algorithm is equal in each benchmark case. The maximum resident memory increases in each benchmark case, but is negligibly small, when compared to standard methods.




Figure 5: Half-open cylindrical geometry in the first three-dimensional benchmark. The geometry is modeled as a single NURBS patch using polynomial degrees $p = \{2, 1, 1\}$ and knot vectors $\Xi_1 = (0, 0, 0, 0.5, 0.5, 1, 1, 1)$, $\Xi_2 = (0, 0, 1, 1)$, $\Xi_3 = (0, 0, 1, 1)$. This case is also studied in [47]


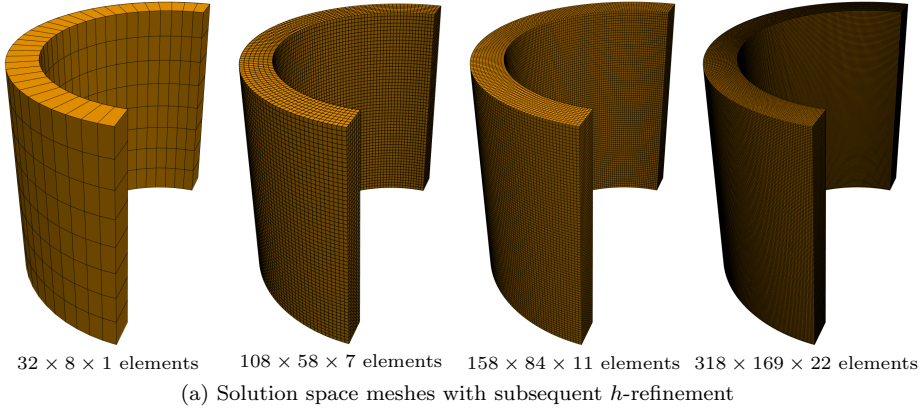


Figure 6: Half-open cylindrical geometry illustrating the solution and interpolation spaces that are used. The meshes from left to right correspond to cases one through four in Table 4. Except for the first case, the mesh size of the interpolation space is roughly half of the mesh size in the solution space. The meshes 2–4 are nearly uniform in each parametric direction $(1, 2, 3)$.




Figure 7: First nine normalized eigenfunctions weighted by the square root of the corresponding eigenvalues in example 1. Results from the benchmark case 4.

Table 4: Enumeration of twenty largest eigenvalues corresponding to the half-open cylinder problem with exponential kernel. The numerical eigenvalues have been computed by the proposed isogeometric Galerkin method employing interpolation based quadrature for the four different cases of solution and interpolation spaces depicted in Figure 6. Computations executed in parallel on 24 threads.

| Mode | Eigenvalue | | | | |
|---|---------------|---------------|----------------|------------------|-------------------------|
| | Case 1 | Case 2 | Case 3 | Case 4 | Benchmark [47] |
| 1 | 162.9780174 | 162.7993245 | 162.7970353 | 162.7964588 | 162.7993688 |
| 2 | 91.60945568 | 91.43092739 | 91.42851673 | 91.42791690 | 91.43079317 |
| 3 | 57.71077642 | 57.56741771 | 57.56494967 | 57.56435726 | 57.56765769 |
| 4 | 51.25928107 | 51.09028325 | 51.08805606 | 51.08750688 | 51.09029418 |
| 5 | 38.94737779 | 38.79796382 | 38.79533658 | 38.79471081 | 38.79808752 |
| 6 | 28.07061519 | 27.90416545 | 27.90174465 | 27.90116381 | 27.90392169 |
| 7 | 25.20395629 | 25.05648777 | 25.05403032 | 25.05344659 | 25.05681566 |
| 8 | 19.47353676 | 19.36965439 | 19.36706688 | 19.36647412 | 19.36866576 |
| 9 | 16.31659674 | 16.15727094 | 16.15488504 | 16.15431116 | 16.15702078 |
| 10 | 15.94684063 | 15.79586713 | 15.79322880 | 15.79261328 | 15.79601336 |
| 11 | 15.25982561 | 15.14713224 | 15.14435936 | 15.14372663 | 15.14613720 |
| 12 | 11.33010899 | 11.21401641 | 11.21139196 | 11.21079288 | 11.21341529 |
| 13 | 10.32857874 | 10.17949575 | 10.17688327 | 10.17627225 | 10.17964708 |
| 14 | 9.850456030 | 9.693951600 | 9.691442730 | 9.690860550 | 9.693535348 |
| 15 | 8.180517780 | 8.055672110 | 8.052867970 | 8.052234280 | 8.055143387 |
| 16 | 7.646028240 | 7.579744590 | 7.577085830 | 7.576500070 | 7.576970969 |
| 17 | 6.864934846 | 6.723174560 | 6.720429030 | 6.719805090 | 6.722768096 |
| 18 | 6.530850643 | 6.447248670 | 6.444416950 | 6.443791080 | 6.444800351 |
| 19 | 6.327439841 | 6.177714250 | 6.175219510 | 6.174649400 | 6.177288031 |
| 20 | 5.893677242 | 5.764614360 | 5.761833040 | 5.761202410 | 5.764324956 |
| Total number of elements ^a | 256 (256) | 43848 (4524) | 145992 (16380) | 1182324 (145992) | 256 (^{-b}) |
| Number of degrees of freedom ^a | 1050 (1050) | 59940 (8525) | 179998 (24948) | 1317384 (179998) | 1050 (^{-b}) |
| Mesh size ^a | 2.857 (2.857) | 0.465 (0.997) | 0.311 (0.645) | 0.155 (0.311) | 2.857 (^{-b}) |
| Mesh size/correlation length ^a | 0.571 (0.571) | 0.093 (0.199) | 0.062 (0.129) | 0.031 (0.062) | 0.571 (^{-b}) |
| Pre-solution formation and assembly time | 0.46 s | 0.59 s | 0.63 s | 1.07 s | ^{-c} |
| Number of iterations | 64 | 63 | 63 | 63 | ^{-c} |
| Maximum resident memory [GB] | 0.347 | 0.376676 | 0.586868 | 2.210916 | ^{-c} |
| Solution time | 0.72 s | 12.91 s | 90.33 s | 1.92 h | > 24 h |

^a [solution space] ([interpolation space])

^b using exact kernel

^c not provided in [47]

The rightmost column summarizes the results obtained by the spectral stochastic isogeometric Galerkin method proposed in [47]. On the same mesh (Case 1) we observe a speed-up of around 5 orders in magnitude. This comparison is not completely fair because full Gaussian quadrature performed in [47] is much more accurate than our interpolation based quadrature technique on the same mesh. It may be observed that qualitatively similar results are obtained with the meshes in Case 2 and 3. The speed-up here is, relatively, four and three orders of magnitude.

6.3. Random field with Gaussian kernel in a three-dimensional hemispherical shell

We consider the hemispherical shell with stiffener depicted in figure 8. This three dimensional shell structure is similar to the model published in [48] but has a slightly different stiffener profile. We use the Gaussian covariance function with a correlation length $bL = 0.5L$, where the characteristic domain length, $L \approx 176$, is the diameter of the stiffener ring.

We perform two case studies across a range of polynomial degrees $p = \{2, 6, 10, 14, 16\}$. The two cases differ in the following way:

- (I) The solution mesh is the same as the interpolation mesh.
- (II) The solution mesh is twice as fine as the interpolation mesh in every component direction.

In both studies we use interpolation and solution meshes obtained by p -refinement of the geometrical model, followed by uniform h -refinement. The continuity of these spaces is thus C^0 at knots that are present in the initial coarse geometrical model and C^{p-1} continuity at new knots introduced by h -refinement.




Figure 8: Hemispherical shell with a stiffener. The geometry is modeled as a single NURBS patch using polynomial degrees $p = \{2, 2, 2\}$ and knot vectors $\Xi_1 = (0, 0, 0, 1, 1, 2, 2, 3, 3, 3)$, $\Xi_2 = (0, 0, 0, 1, 1, 2, 2, 3, 3, 4, 4, 4)$, $\Xi_3 = (0, 0, 0, 1, 1, 1)$.

Case I

Figures 9a and 9b depict the interpolation and solution meshes, respectively, that are used in this benchmark. As described they are identical. The numerical results are summarized in Table 5. As in the previous three-dimensional benchmark problem the presolution setup costs hardly contribute to the total cost of the method. Again, the main computational cost lies in the matrix-free matrix-vector product that is evaluated in each iteration of the Lanczos eigenvalue solver. Due to the increased number of degrees of freedom for higher polynomial degrees the associated computational cost increases. Interestingly, the number of iterations is reduced in the case $p = 16$ as compared to the lower polynomial degrees.

Table 5: Enumeration of twenty largest eigenvalues corresponding to the hemispherical shell with stiffener problem with Gaussian kernel. The numerical eigenvalues have been computed by the proposed isogeometric Galerkin method employing interpolation based quadrature using the interpolation and solution mesh depicted in Figure 9a and 9b, respectively. Computations are executed in parallel on 12 threads.



| Mode | Eigenvalue | | | | |
|--|---------------|---------------|---------------|-----------------|-----------------|
| | $p = 2$ | $p = 6$ | $p = 10$ | $p = 14$ | $p = 16$ |
| 1 | 14548.14836 | 14502.90405 | 14476.37478 | 14476.28290 | 14476.27783 |
| 2 | 6590.886122 | 6553.622930 | 6531.848385 | 6531.774420 | 6531.770729 |
| 3 | 6590.886122 | 6553.622930 | 6531.848385 | 6531.774420 | 6531.770728 |
| 4 | 2118.360157 | 2101.642900 | 2091.911569 | 2091.878625 | 2091.877012 |
| 5 | 2118.358449 | 2101.642996 | 2091.911569 | 2091.878625 | 2091.877003 |
| 6 | 1994.090912 | 1979.961757 | 1971.740772 | 1971.739918 | 1971.747135 |
| 7 | 558.7902107 | 554.9316152 | 552.6952326 | 552.7038312 | 552.7086269 |
| 8 | 558.7902107 | 554.9316152 | 552.6952326 | 552.7038312 | 552.7086185 |
| 9 | 530.3814344 | 525.2418684 | 522.2580040 | 522.2479339 | 522.2474594 |
| 10 | 530.3814344 | 525.2418684 | 522.2580040 | 522.2479339 | 522.2474525 |
| 11 | 114.2741328 | 113.5448612 | 113.1241037 | 113.1279916 | 113.1296326 |
| 12 | 114.2740407 | 113.5448612 | 113.1241037 | 113.1279916 | 113.1296295 |
| 13 | 107.9782997 | 106.7763105 | 106.0789732 | 106.0766531 | 106.0765638 |
| 14 | 107.9781271 | 106.7763105 | 106.0789732 | 106.0766530 | 106.0765540 |
| 15 | 101.9364272 | 101.4389135 | 101.1503447 | 101.1560706 | 101.1581902 |
| 16 | 21.52017977 | 21.44679212 | 21.40372683 | 21.40541424 | 21.40600231 |
| 17 | 21.52017977 | 21.44679212 | 21.40372683 | 21.40541424 | 21.40599155 |
| 18 | 19.28909583 | 19.17852048 | 19.11482729 | 19.11584295 | 19.11622952 |
| 19 | 19.28909583 | 19.17852048 | 19.11482729 | 19.11584292 | 19.11622323 |
| 20 | 18.44129352 | 18.21491774 | 18.08322462 | 18.08280013 | 18.08278502 |
| Number of elements ^a | 3864 (3864) | 3864 (3864) | 3864 (3864) | 3864 (3864) | 3864 (3864) |
| Number of degrees of freedom ^a | 9612 (9612) | 32760 (32760) | 74052 (74052) | 138096 (138096) | 180090 (180090) |
| Max. mesh size ^a | 7.992 (7.992) | 7.992 (7.992) | 7.992 (7.992) | 7.992 (7.992) | 7.992 (7.992) |
| Max. mesh size/correlation length ^a | 0.091 (0.091) | 0.091 (0.091) | 0.091 (0.091) | 0.091 (0.091) | 0.091 (0.091) |
| Pre-solution formation and assembly time | 0.51 s | 0.58 s | 0.70 s | 0.81 s | 0.88 s |
| Number of iterations | 52 | 52 | 52 | 52 | 41 |
| Maximum resident memory [GB] | 0.51 | 0.577 | 0.702 | 1.028 | 1.339 |
| Solution time | 15.28 s | 160.03 s | 0.210 h | 1.175 h | 1.65 h |

^a[solution space] ([interpolation space])



Case II

In the second benchmark the element size of the solution mesh is halved, see Figure 9c. The interpolation mesh is kept the same as in the first benchmark. The obtained results are presented in Table 6. By comparing these results with those in Table 5 it may be observed that the dimension of the solution space does not significantly affect the total solver costs. Indeed, the number of degrees of freedom are more than doubled, yet the timings stay more or less the same, independent of polynomial degree. The increased dimension of the solution space mesh is reflected in the maximum resident memory, which has increased as compared to the results in Table 5. As witnessed in the previous benchmark, the higher order simulations required fewer iterations than the lower order ones.



Remark 6.3. *Note that the flexibility in mesh size of the interpolation versus the trial space mesh provides a mechanism by which the error due to quadrature versus the error due to discretization can be effectively controlled.*



(a) Interpolation space mesh in case I and II with $23 \times 84 \times 2$ elements in parametric directions $(1, 2, 3)$



(b) Solution space mesh in case I with $23 \times 84 \times 2$ elements in parametric directions $(1, 2, 3)$



(c) Solution space mesh in case II with $46 \times 168 \times 4$ elements in parametric directions $(1, 2, 3)$

Figure 9: The solution space and interpolation space meshes and the cross-sections used in example 2a and 2b.

Table 6: Enumeration of twenty largest eigenvalues corresponding to the hemispherical shell with stiffener problem with Gaussian kernel. The numerical eigenvalues have been computed by the proposed isogeometric Galerkin method employing interpolation based quadrature using the interpolation and solution mesh depicted in Figure 9a and 9c, respectively. Computations are executed in parallel on 12 threads.

| Mode | Eigenvalue | | | | |
|--|---------------|----------------|----------------|-----------------|-----------------|
| | $p = 2$ | $p = 6$ | $p = 10$ | $p = 14$ | $p = 16$ |
| 1 | 14548.14836 | 14502.90404 | 14476.37478 | 14476.28290 | 14476.27783 |
| 2 | 6590.886122 | 6553.622930 | 6531.848385 | 6531.774420 | 6531.770729 |
| 3 | 6590.886122 | 6553.622930 | 6531.848385 | 6531.774420 | 6531.770725 |
| 4 | 2118.360157 | 2101.642996 | 2091.911569 | 2091.878625 | 2091.877015 |
| 5 | 2118.358449 | 2101.642996 | 2091.911569 | 2091.878625 | 2091.877012 |
| 6 | 1994.090912 | 1979.961757 | 1971.740772 | 1971.739918 | 1971.747136 |
| 7 | 558.7902107 | 554.9316152 | 552.6952326 | 552.7038312 | 552.7086262 |
| 8 | 558.7902107 | 554.9316152 | 552.6952326 | 552.7038312 | 552.7086255 |
| 9 | 530.3814344 | 525.2418684 | 522.2580040 | 522.2479339 | 522.2474587 |
| 10 | 530.3814344 | 525.2418684 | 522.2580040 | 522.2479339 | 522.2474574 |
| 11 | 114.2741328 | 113.5448612 | 113.1241037 | 113.1279916 | 113.1296371 |
| 12 | 114.2740407 | 113.5448612 | 113.1241037 | 113.1279916 | 113.1296317 |
| 13 | 107.9782997 | 106.7763105 | 106.0789732 | 106.0766530 | 106.0765590 |
| 14 | 107.9781271 | 106.7763105 | 106.0789732 | 106.0766530 | 106.0765573 |
| 15 | 101.9364271 | 101.4389135 | 101.1503447 | 101.1560706 | 101.1581809 |
| 16 | 21.52017977 | 21.44679212 | 21.40372683 | 21.40541425 | 21.40600074 |
| 17 | 21.52017977 | 21.44679212 | 21.40372683 | 21.40541424 | 21.40599922 |
| 18 | 19.28909583 | 19.17852048 | 19.11482729 | 19.11584292 | 19.11623197 |
| 19 | 19.28909583 | 19.17852048 | 19.11482729 | 19.11584292 | 19.11623140 |
| 20 | 18.44129352 | 18.21491774 | 18.08322462 | 18.08280013 | 18.08278785 |
| Number of elements ^a | 30912 (3864) | 30912 (3864) | 30912 (3864) | 30912 (3864) | 30912 (3864) |
| Number of degrees of freedom ^a | 51900 (9612) | 117180 (32760) | 212380 (74052) | 342108 (138096) | 421360 (180090) |
| Max. mesh size ^a | 4.019 (7.992) | 4.019 (7.992) | 4.019 (7.992) | 4.019 (7.992) | 4.019 (7.992) |
| Max. mesh size/correlation length ^a | 0.046 (0.091) | 0.046 (0.091) | 0.046 (0.091) | 0.046 (0.091) | 0.046 (0.091) |
| Pre-solution formation and assembly time | 0.63 s | 0.67 s | 0.74 s | 0.90 s | 1.00 s |
| Number of iterations | 52 | 52 | 52 | 41 | 41 |
| Maximum resident memory [GB] | 0.546 | 0.891 | 2.290 | 5.276 | 7.483 |
| Solution time | 16.86 s | 162.78 s | 0.256 h | 0.912 h | 1.63 h |

^a[solution space] ([interpolation space])

7. Conclusion

This paper presented an efficient matrix-free Galerkin method for the Karhunen-Loève (KL) series expansion of random fields. The KL expansion requires the solution of a generalized eigenvalue problem corresponding to the homogeneous Fredholm integral eigenvalue problem of the second kind, and is computationally challenging for several reasons. Firstly, the Galerkin method requires numerical integration over a $2d$ -dimensional domain, where d , in this work, denotes the spatial dimension. Consequently, classical formation and assembly procedures have a time complexity that scales $\mathcal{O}(N_e^2 \cdot p^{3d})$ with increasing polynomial degree p and number of elements N_e . Secondly, the main system matrix is dense and requires $\mathcal{O}(N^2)$ bytes of storage, where N is the global number of degrees of freedom. This means that a discretization involving a hundred thousand degrees of freedom requires at least 80GB of RAM to store the main system matrix in double precision. Hence, the computational complexity as well as memory requirements of standard solution techniques become quickly computationally intractable with increasing polynomial degree, problem size and spatial dimension.

We proposed an efficient solution methodology that significantly ameliorates the aforementioned computational challenges. Our approach is based on the following key ingredients

1. A trial space of rational spline functions, whose Gramian or mass matrix has a Kronecker structure independent of the geometric mapping;
2. An inexpensive reformulation of the generalized algebraic eigenvalue problem into a standard algebraic eigenvalue problem;
3. A degenerate kernel approximation of the covariance function using smooth tensor product splines;

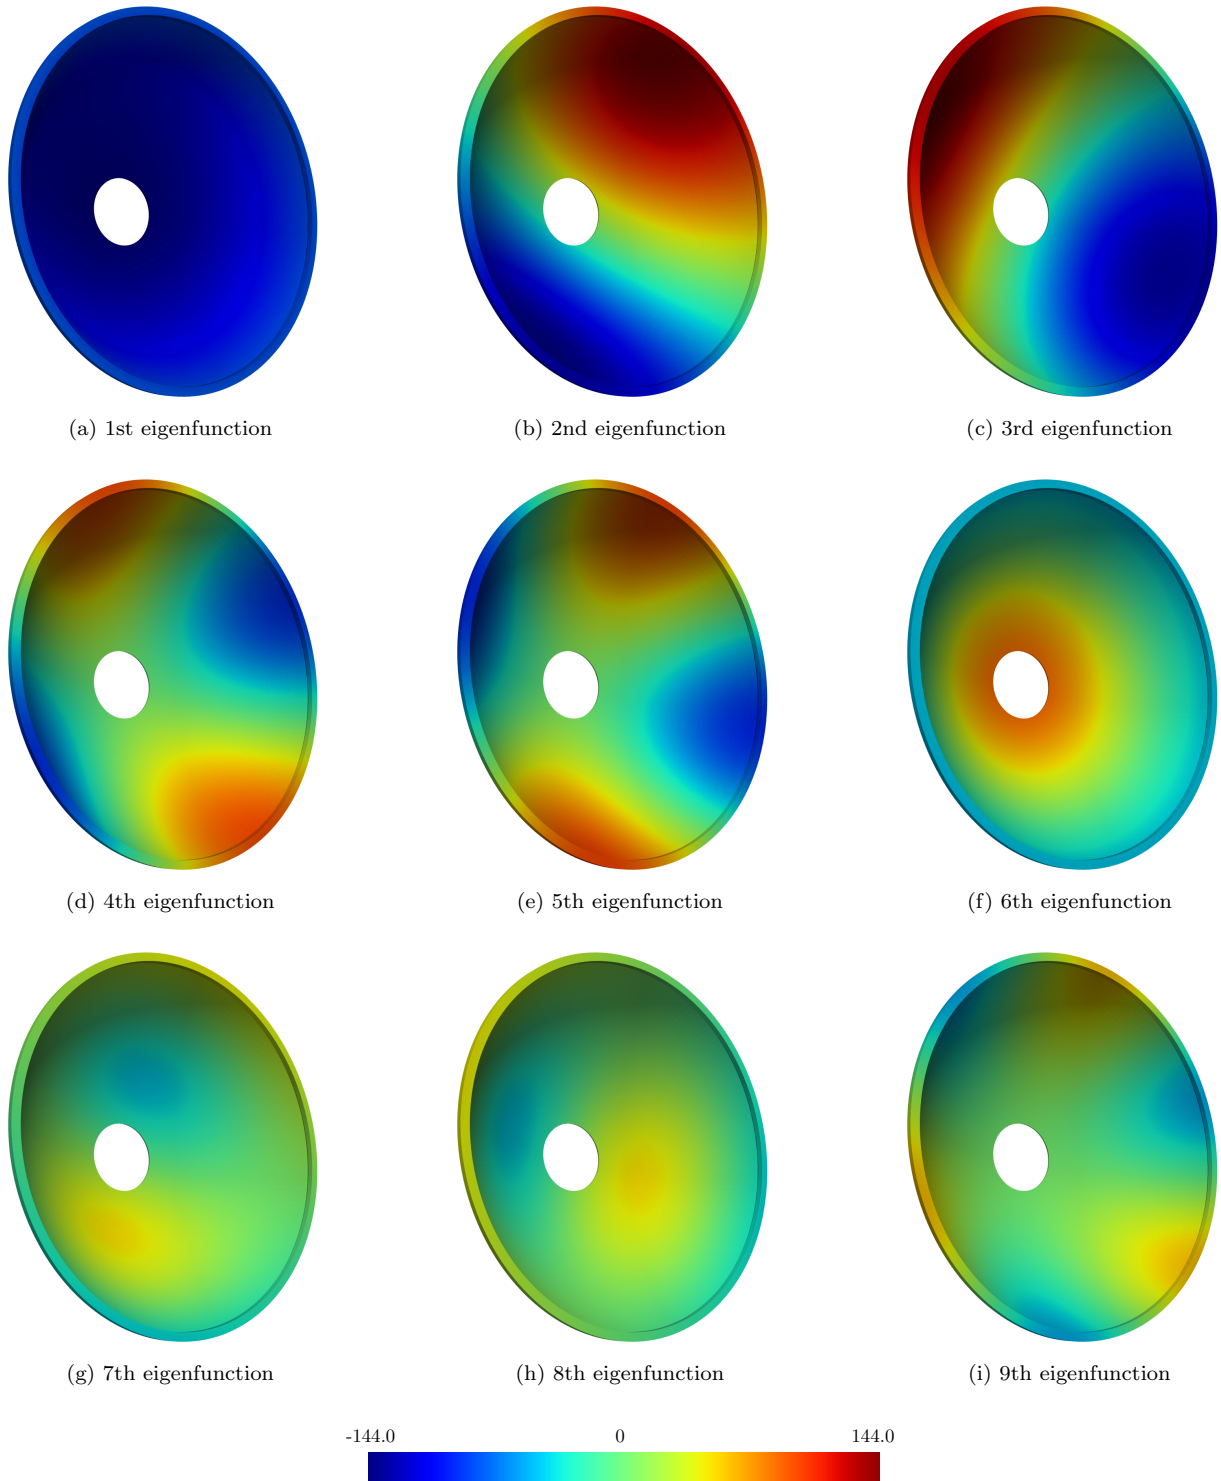


Figure 10: First nine normalized eigenfunctions weighted by the square root of the corresponding eigenvalues in example 2b

4. Formulation of an efficient matrix-free and multi-threaded matrix-vector product for iterative eigenvalue solvers, which utilizes the Kronecker structure of the system matrices.

In Step 2 the reformulation to a standard eigenvalue problem significantly reduces the computational cost while improving conditioning. This can be done efficiently due to the Kronecker structure of the mass matrix, which is a result of the particular choice of trial space, see Step 1. In Step 3 the degenerate kernel approximation enables us to evaluate the resulting integrals exactly with a minimal number of evaluation points. Both steps involve matrices that are endowed with a Kronecker structure and can be performed matrix-free in $\mathcal{O}(N \cdot N^{1/d})$ time. The leading cost of the method is due to the Lanczos eigenvalue algorithm, which involves dense matrix-vector multiplications. As noted in Step 4, we perform this step matrix-free, by computing the necessary components on the fly, and in parallel in approximately $\mathcal{O}(N^2 N_{\text{iter}}/N_{\text{threads}})$ time. Here N_{iter} denotes the number of iterations of the eigensolver and N_{threads} is the number of simultaneous processes. Several three dimensional benchmark problems involving non-trivial geometrical mappings have illustrated exceptional efficiency and effectiveness of the proposed solution methodology. In particular, we showed that the proposed methodology scales favorably with polynomial degree and works particularly well for smooth covariance functions, such as the Gaussian kernel. The Python implementation used to generate these results has been provided as open-source software and is available for download at <https://github.com/mika05/tensiga>.

In a follow-up study we plan to extensively study the accuracy of the proposed solution methodology. There are two sources of error: (1) a quadrature error due to the degenerate kernel approximation of the covariance function; and (2) a discretization error due to the finite dimensional representation of the eigenmodes. We will perform a-priori as well as a-posteriori error analysis and formulate criteria for bounding the error due to quadrature by the discretization error. Within the same context of accuracy and robustness it is interesting to extend the spectral analysis results in [23] to generalized eigenvalue problems corresponding to Fredholm integral equations for different covariance functions as well as polynomial order of approximation.

We also plan to further improve the efficiency of the proposed method where possible. The proposed matrix-free algorithm lends itself for acceleration on graphics processing units (GPU's). Furthermore, exploiting particular structure (such as sparsity or symmetry) of the covariance function may lead to improved solver cost. For example, the hierarchical matrix method proposed in [29] performs the matrix-vector products in $\mathcal{O}(N \log N)$ time, by exploiting certain structure of the covariance function.

There are several other interesting avenues for future research. Some or all of the techniques proposed here could be applied to linear as well as non-linear Fredholm integral differential equations of the first as well as the second kind. While the proposed method is designed for smooth kernels it would be interesting to develop similar methods that are tailored towards continuous kernels, such as the exponential kernel, or even singular kernels, which are typical in boundary integral equations.

Finally we would like to mention that similar techniques can be applied in the context of the collocation method. The computational cost of such a method would be similar to that of the proposed Galerkin method.

Acknowledgments

M.L. Mika, R.R. Hiemstra and D. Schillinger gratefully acknowledge funding from the German Research Foundation through the DFG Emmy Noether Award SCH 1249/2-1. T.J.R. Hughes and R.R. Hiemstra were partially supported by the National Science Foundation Industry/University Cooperative Research Center (IUCRC) for Efficient Vehicles and Sustainable Transportation Systems (EV-STs), and the United States Army CCDC Ground Vehicle Systems Center (TARDEC/NSF Project # 1650483 AMD 2). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. The authors thank Mona Dannert and Udo Nackenhorst for very helpful discussions and comments.

References

- [1] ALLAIX, D. L., AND CARBONE, V. I. Karhunen-Loève decomposition of random fields based on a hierarchical matrix approach. *International Journal for Numerical Methods in Engineering* 94, 11 (June 2013), 1015–1036.
- [2] ARTHUR, D. W. The Solution of Fredholm Integral Equations Using Spline Functions. *IMA Journal of Applied Mathematics* 11, 2 (1973), 121–129.
- [3] ATKINSON, K. E. *The Numerical Solution of Integral Equations of the Second Kind*, 1 ed. Cambridge University Press, June 1997.
- [4] BETZ, W., PAPAIOANNOU, I., AND STRAUB, D. Numerical methods for the discretization of random fields by means of the Karhunen-Loève expansion. *Computer Methods in Applied Mechanics and Engineering* 271 (Apr. 2014), 109–129.
- [5] CALABRÒ, F., SANGALLI, G., AND TANI, M. Fast formation of isogeometric Galerkin matrices by weighted quadrature. *Computer Methods in Applied Mechanics and Engineering* 316 (Apr. 2017), 606–622.
- [6] COTTRELL, J., HUGHES, T., AND REALI, A. Studies of refinement and continuity in isogeometric structural analysis. *Computer Methods in Applied Mechanics and Engineering* 196, 41-44 (Sept. 2007), 4160–4183.
- [7] COTTRELL, J. A., HUGHES, T. J. R., AND BAZILEVS, Y. *Isogeometric analysis: toward integration of CAD and FEA*. Wiley, Chichester, West Sussex, U.K. ; Hoboken, NJ, 2009. OCLC: ocn335682757.
- [8] COURANT, R., AND HILBERT, D. *Methods of Mathematical Physics*, 1 ed. Wiley, Apr. 1989.
- [9] DAHMEN, W., HARBRECHT, H., AND SCHNEIDER, R. Compression Techniques for Boundary Integral Equations—Asymptotically Optimal Complexity Estimates. *SIAM Journal on Numerical Analysis* 43, 6 (Jan. 2006), 2251–2271.
- [10] DE BOOR, C. *A practical guide to splines*. No. 27 in Applied mathematical sciences. Springer-Verlag, New York, 1978.
- [11] DING, C., HU, X., CUI, X., LI, G., CAI, Y., AND TAMMA, K. K. Isogeometric generalized n th order perturbation-based stochastic method for exact geometric modeling of (composite) structures: Static and dynamic analysis with random material parameters. *Computer Methods in Applied Mechanics and Engineering* 346 (Apr. 2019), 1002–1024.
- [12] ECKERT, C., BEER, M., AND SPANOS, P. D. A polynomial chaos method for arbitrary random inputs using B-splines. *Probabilistic Engineering Mechanics* 60 (Apr. 2020), 103051.
- [13] EIERMANN, M., ERNST, O. G., AND ULLMANN, E. Computational aspects of the stochastic finite element method. *Computing and Visualization in Science* 10, 1 (Feb. 2007), 3–15.
- [14] GHANEM, R. G., AND SPANOS, P. D. *Stochastic Finite Elements: A Spectral Approach*. Springer New York, New York, NY, 1991.
- [15] GOLUB, G. H., AND VAN LOAN, C. F. *Matrix computations*, 3rd ed ed. Johns Hopkins studies in the mathematical sciences. Johns Hopkins University Press, Baltimore, 1996.
- [16] HACKBUSCH, W., KHOROMSKIJ, B. N., AND TYRTYSHNIKOV, E. E. Hierarchical Kronecker tensor-product approximations. *Journal of Numerical Mathematics* 13, 2 (Jan. 2005).
- [17] HARBRECHT, H., PETERS, M., AND SIEBENMORGEN, M. Efficient approximation of random fields for numerical applications. *Numerical Linear Algebra with Applications* 22, 4 (Aug. 2015), 596–617.
- [18] HIEMSTRA, R. R., CALABRÒ, F., SCHILLINGER, D., AND HUGHES, T. J. Optimal and reduced quadrature rules for tensor product and hierarchically refined splines in isogeometric analysis. *Computer Methods in Applied Mechanics and Engineering* 316 (2017), 966 – 1004. Special Issue on Isogeometric Analysis: Progress and Challenges.
- [19] HIEMSTRA, R. R., SANGALLI, G., TANI, M., CALABRÒ, F., AND HUGHES, T. J. Fast formation and assembly of finite element matrices with application to isogeometric linear elasticity. *Computer Methods in Applied Mechanics and Engineering* 355 (Oct. 2019), 234–260.
- [20] HIEN, T. D., AND NOH, H.-C. Stochastic isogeometric analysis of free vibration of functionally graded plates considering material randomness. *Computer Methods in Applied Mechanics and Engineering* 318 (May 2017), 845–863.
- [21] HUGHES, T., COTTRELL, J., AND BAZILEVS, Y. Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. *Computer Methods in Applied Mechanics and Engineering* 194, 39-41 (Oct. 2005), 4135–4195.
- [22] HUGHES, T., REALI, A., AND SANGALLI, G. Duality and unified analysis of discrete approximations in structural dynamics and wave propagation: Comparison of p-method finite elements with k-method NURBS. *Computer Methods in Applied Mechanics and Engineering* 197, 49-50 (Sept. 2008), 4104–4124.
- [23] HUGHES, T. J., EVANS, J. A., AND REALI, A. Finite element and NURBS approximations of eigenvalue, boundary-value, and initial-value problems. *Computer Methods in Applied Mechanics and Engineering* 272 (Apr. 2014), 290–320.
- [24] JAHANBIN, R., AND RAHMAN, S. An isogeometric collocation method for efficient random field discretization. *International Journal for Numerical Methods in Engineering* 117, 3 (Jan. 2019), 344–369.
- [25] JAHANBIN, R., AND RAHMAN, S. Stochastic isogeometric analysis in linear elasticity. *Computer Methods in Applied Mechanics and Engineering* 364 (June 2020), 112928.
- [26] KARHUNEN, K. *Über lineare Methoden in der Wahrscheinlichkeitsrechnung*. Suomalaisen Tiedeakatemian toimituksia. Zugl.: Helsinki, Univ., Diss., 1947, Helsinki, 1947.
- [27] KEESE, A. A Review of Recent Developments in the Numerical Solution of Stochastic Partial Differential Equations (Stochastic Finite Elements). *Braunschweig, Institut für Wissenschaftliches Rechnen* (2003).
- [28] KEESE, A. Numerical Solution of Systems with Stochastic Uncertainties: A General Purpose Framework for Stochastic Finite Elements, June 2004.
- [29] KHOROMSKIJ, B. N., LITVINENKO, A., AND MATTHIES, H. G. Application of hierarchical matrices for computing the Karhunen-Loève expansion. *Computing* 84, 1-2 (Apr. 2009), 49–67.
- [30] KJOLSTAD, F., KAMIL, S., CHOU, S., LUGATO, D., AND AMARASINGHE, S. The tensor algebra compiler. *Proceedings of the ACM on Programming Languages* 1, OOPSLA (Oct. 2017), 1–29.

- [31] KRESS, R. *Linear integral equations*, third edition ed. No. volume 82 in Applied mathematical sciences. Springer, New York, 2014.
- [32] LAM, S. K., PITROU, A., AND SEIBERT, S. Numba: a LLVM-based Python JIT compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC - LLVM '15* (Austin, Texas, 2015), ACM Press, pp. 1–6.
- [33] LATTNER, C., AND ADVE, V. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-Directed and Runtime Optimization* (USA, 2004), CGO '04, IEEE Computer Society, p. 75. event-place: Palo Alto, California.
- [34] LI, K., GAO, W., WU, D., SONG, C., AND CHEN, T. Spectral stochastic isogeometric analysis of linear elasticity. *Computer Methods in Applied Mechanics and Engineering* 332 (Apr. 2018), 157–190.
- [35] LI, K., WU, D., AND GAO, W. Spectral stochastic isogeometric analysis for static response of FGM plate with material uncertainty. *Thin-Walled Structures* 132 (Nov. 2018), 504–521.
- [36] LI, K., WU, D., AND GAO, W. Spectral stochastic isogeometric analysis for linear stability analysis of plate. *Computer Methods in Applied Mechanics and Engineering* 352 (Aug. 2019), 1–31.
- [37] LI, K., WU, D., GAO, W., AND SONG, C. Spectral stochastic isogeometric analysis of free vibration. *Computer Methods in Applied Mechanics and Engineering* 350 (June 2019), 1–27.
- [38] LOÈVE, M. Fonctions aleatoires du second ordre. Processus stochastique et mouvement Brownien. *Paris, Gauthier-Villars* (1948), 366–420.
- [39] MANTZAFARIS, A., AND JÜTTLER, B. Integration by interpolation and look-up for Galerkin-based isogeometric analysis. *Computer Methods in Applied Mechanics and Engineering* 284 (Feb. 2015), 373–400.
- [40] MELCHERS, R. E., AND BECK, A. T., Eds. *Structural Reliability Analysis and Prediction*. John Wiley & Sons Ltd, Chichester, UK, Oct. 2017.
- [41] MERCER, J. Functions of Positive and Negative Type, and their Connection with the Theory of Integral Equations. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 209, 441–458 (Jan. 1909), 415–446.
- [42] ORSZAG, S. A. Spectral methods for problems in complex geometries. *Journal of Computational Physics* 37, 1 (Aug. 1980), 70–92.
- [43] PAN, M., JÜTTLER, B., AND GIUST, A. Fast formation of isogeometric Galerkin matrices via integration by interpolation and look-up. *Computer Methods in Applied Mechanics and Engineering* 366 (July 2020), 113005.
- [44] PHOON, K., HUANG, S., AND QUEK, S. Implementation of Karhunen–Loève expansion for simulation using a wavelet-Galerkin scheme. *Probabilistic Engineering Mechanics* 17, 3 (July 2002), 293–303.
- [45] PIEGL, L., AND TILLER, W. *The NURBS Book*. Monographs in Visual Communications. Springer Berlin Heidelberg, Berlin, Heidelberg, 1995.
- [46] PUZYREV, V., DENG, Q., AND CALO, V. Spectral approximation properties of isogeometric analysis with variable continuity. *Computer Methods in Applied Mechanics and Engineering* 334 (June 2018), 22–39.
- [47] RAHMAN, S. A Galerkin isogeometric method for Karhunen–Loève approximation of random fields. *Computer Methods in Applied Mechanics and Engineering* 338 (Aug. 2018), 533–561.
- [48] RANK, E., DÜSTER, A., NÜBEL, V., PREUSCH, K., AND BRUHNS, O. High order finite elements for shells. *Computer Methods in Applied Mechanics and Engineering* 194, 21–24 (June 2005), 2494–2512.
- [49] SAAD, Y. *Numerical methods for large eigenvalue problems*, rev. ed ed. No. 66 in Classics in applied mathematics. Society for Industrial and Applied Mathematics, Philadelphia, 2011.
- [50] SANDE, E., MANNI, C., AND SPELEERS, H. Explicit error estimates for spline approximation of arbitrary smoothness in isogeometric analysis. *Numerische Mathematik* 144, 4 (Apr. 2020), 889–929.
- [51] SANGALLI, G., AND TANI, M. Matrix-free weighted quadrature for a computationally efficient isogeometric k -method. *Computer Methods in Applied Mechanics and Engineering* 338 (Aug. 2018), 117–133.
- [52] SCHWAB, C., AND TODOR, R. A. Karhunen–Loève approximation of random fields by generalized fast multipole methods. *Journal of Computational Physics* 217, 1 (Sept. 2006), 100–122.
- [53] SCIPY 1.0 CONTRIBUTORS, VIRTANEN, P., GOMMERS, R., OLIPHANT, T. E., HABERLAND, M., REDDY, T., COURNAPEAU, D., BUROVSKI, E., PETERSON, P., WECKESSER, W., BRIGHT, J., VAN DER WALT, S. J., BRETT, M., WILSON, J., MILLMAN, K. J., MAYOROV, N., NELSON, A. R. J., JONES, E., KERN, R., LARSON, E., CAREY, C. J., POLAT, I., FENG, Y., MOORE, E. W., VANDERPLAS, J., LAXALDE, D., PERKTOLD, J., CIRMANN, R., HENRIKSEN, I., QUINTERO, E. A., HARRIS, C. R., ARCHIBALD, A. M., RIBEIRO, A. H., PEDREGOSA, F., AND VAN MULBREGT, P. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods* 17, 3 (Mar. 2020), 261–272.
- [54] SUDRET, B., AND KUYREGHIAN, A. *Stochastic finite element methods and reliability: a state-of-the-art report*. Berkeley, Department of Civil and Environmental Engineering, University of California, 2000.
- [55] TINO EIBNER, J. M. M. Fast algorithms for setting up the stiffness matrix in hp-FEM: a comparison, 2005.
- [56] VAN DER WALT, S., COLBERT, S. C., AND VAROQUAUX, G. The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science & Engineering* 13, 2 (Mar. 2011), 22–30.
- [57] VOS, P. E., SHERWIN, S. J., AND KIRBY, R. M. From h to p efficiently: Implementing finite and spectral/hp element methods to achieve optimal performance for low- and high-order discretisations. *Journal of Computational Physics* 229, 13 (July 2010), 5161–5181.
- [58] WANG, W., CHEN, G., YANG, D., AND KANG, Z. Stochastic isogeometric analysis method for plate structures with random uncertainty. *Computer Aided Geometric Design* 74 (Oct. 2019), 101772.
- [59] XIU, D., AND KARNIADAKIS, G. E. The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations. *SIAM Journal on Scientific Computing* 24, 2 (Jan. 2002), 619–644.
- [60] ZHANG, H., AND SHIBUTANI, T. Development of stochastic isogeometric analysis (SIGA) method for uncertainty in shape.

International Journal for Numerical Methods in Engineering (Dec. 2018), nme.6008.